arXiv:1510.00697v1 [cs.PL] 2 Oct 2015

Towards a Fully Abstract Compiler Using Micro-Policies

— Secure Compilation for Mutually Distrustful Components —

Yannis Juglaret[1,2] Cătălin Hriţcu[1] Arthur Azevedo de Amorim[3]

Benjamin C. Pierce[3] Antal Spector-Zabusky[3] Andrew Tolmach[4]

[1]Inria Paris [2]Université Paris Diderot (Paris 7)

[3]University of Pennsylvania [4]Portland State University

Technical Report

initial version: August 21, 2015

last revised: October 5, 2015

Abstract

Secure compilation prevents all low-level attacks on compiled code and allows for sound reasoning about security in the source language. In this work we propose a new attacker model for secure compilation that extends the well-known notion of full abstraction to ensure protection for mutually distrustful components. We devise a compiler chain (compiler, linker, and loader) and a novel security monitor that together defend against this strong attacker model. The monitor is implemented using a recently proposed, generic tag-based protection framework called micro-policies, which comes with hardware support for efficient caching and with a formal verification methodology. Our monitor protects the abstractions of a simple object-oriented language—class isolation, the method call discipline, and type safety—against arbitrary low-level attackers.

Contents

1 Introduction **4**
- 1.1 General Context . . . 4
- 1.2 Research Problem . . . 4
- 1.3 Our Contribution . . . 5
- 1.4 Other Insights . . . 5

2 Stronger Attacker Model for Secure Compilation of Mutually Distrustful Components **5**
- 2.1 Full Abstraction . . . 5
- 2.2 Limitations of Full Abstraction . . . 6
- 2.3 Mutual Distrust Attacker Model . . . 7

3 Micro-Policies and the PUMP: Efficient Tag-Based Security Monitors **8**

4 Compilation Chain Overview **9**

5 Languages and Machines **10**
- 5.1 Source Level: An Object-Oriented Language . . . 10
  - 5.1.1 Interfacing: A Specification Language for Communicating Classes . . . 10
  - 5.1.2 Source Syntax and Semantics . . . 11
- 5.2 Intermediate Level: An Object-Oriented Stack Machine . . . 12
- 5.3 Target Level: An Extended Symbolic Micro-Policy Machine . . . 13
  - 5.3.1 Symbolic Micro-Policy Machine . . . 13
  - 5.3.2 Extensions to Monitoring Mechanism . . . 14
  - 5.3.3 Easing Component-Oriented Reasoning: Segmented Memory . . . 14

6 Our Solution: Protecting Compiled Code with a Micro-Policy **15**
- 6.1 Two-Step Compilation . . . 15
  - 6.1.1 From Source to Intermediate . . . 15
  - 6.1.2 From Intermediate to Target . . . 16
- 6.2 Micro-Policy Protecting Abstractions . . . 19
  - 6.2.1 Enforcing Class Isolation via Compartmentalization . . . 19
  - 6.2.2 Enforcing Method Call Discipline using Method Entry Points, Linear Return Capabilities, and Register Cleaning . . . 19
  - 6.2.3 Enforcing Type Safety Dynamically . . . 20
  - 6.2.4 Micro-Policy in Detail . . . 21

7 Related Work **22**
- 7.1 Secure Compilation . . . 22
- 7.2 Multi-Language Approaches . . . 23

8 Discussion and Future Work **23**
- 8.1 Finite Memory and Full Abstraction . . . 23
- 8.2 Efficiency and Transparency . . . 24
- 8.3 Future Work . . . 24
- 8.4 Scaling to Real-World Languages . . . 25

A Appendix **25**
- A.1 Encoding Usual Types . . . 25
- A.2 Source Semantics . . . 25

1 Introduction

1.1 General Context

In this work we study compiled partial programs evolving within a low-level environment, with which they can interact. Such interaction is useful — think of high-level programs performing low-level library calls, or of a browser interacting with native code that was sent over the internet [12, 45] — but also dangerous: parts of the environment could be malicious or compromised and try to compromise the program as well [12, 15, 45]. Low-level libraries written in C or in C++ can be vulnerable to control hijacking attacks [15, 41] and be taken over by a remote attacker. When the environment can't be trusted, it is a major concern to ensure the security of running programs.

With today's compilers, low-level attackers [15] can circumvent high-level abstractions [1, 25] and are thus much more powerful than high-level attackers, which means that the security reasoning has to be done at the lowest level, which it is extremely difficult. An alternative is to build a *secure compiler* that ensures that low- and high-level attackers have exactly the same power, allowing for easier, source-level security reasoning [4, 19, 22, 32]. Formally, the notion of secure compilation is usually expressed as *full abstraction* of the translation [1]. Full abstraction is a much stronger property than just compiler correctness [27].

Secure compilation is, however, very hard to achieve in practice. Efficiency, which is crucial for broad adoption [41], is the main challenge. Another concern is transparency. While we want to constrain the power of low-level attackers, the constraints we set should be relaxed enough that there is a way for all benign low-level environments to respect them. If we are not transparent enough, the partial program might be prevented from properly interacting with its environment (e.g. the low-level libraries it requires).

For a compiler targeting machine code, which lacks structure and checks, a typical low-level attacker has write access to the whole memory, and can redirect control flow to any location in memory [15]. Techniques have been developed to deal with such powerful attackers, in particular involving randomization [4] and binary code rewriting [16, 29]. The first ones only offer weak probabilistic guarantees; as a consequence, address space layout randomization [4] is routinely circumvented in practical attacks [17, 37]. The second ones add extra software checks which often come at a high performance cost.

Using additional protection in the hardware can result in secure compilers with strong guarantees [32], without sacrificing efficiency or transparency. Because updating hardware is expensive and hardware adoption takes decades, the need for generic protection mechanisms that can fit with ever-evolving security requirements has emerged. Active research in the domain includes *capability machines* [11, 42, 43] and *tag-based architectures* [8, 9, 14, 39]. In this work, we use a generic tag-based protection mechanism called *micro-policies* [9, 14] as the target of a secure compiler.

Micro-policies provide instruction-level monitoring based on fine-grained metadata tags. In a micro-policy machine, every word of data is augmented with a word-sized tag, and a hardware-accelerated monitor propagates these tags every time a machine instruction gets executed. Micro-policies can be described as a combination of software-defined rules and monitor services. The rules define how the monitor will perform tag propagation instruction-wise, while the services allow for direct interaction between the running code and the monitor. This mechanism comes with an efficient hardware implementation built on top of a RISC processor [14] as well as a mechanized metatheory [9], and has already been used to enforce a variety of security policies [9, 14].

1.2 Research Problem

Recent work [6, 32] has illustrated how *protected module architectures* — a class of hardware architectures featuring coarse-grained isolation mechanisms [20, 28, 38] — can help in devising a fully abstract compilation scheme for a Java-like language. This scheme assumes the compiler knows which components in the program can be trusted and which ones cannot, and protects the trusted components from the distrusted ones by isolating them in a protected module.

This kind of protection is only appropriate when all the components we want to protect can be trusted, for example because they have been verified [5]. Accounting for the cases in which this is not possible, we present and adopt a stronger attacker model of *mutual distrust*: in this setting a secure compiler should protect each component from every other component, so that whatever the compromise scenario may be, uncompromised components always get protected from the compromised ones.

The main questions we address in this work are: (1) can we build a fully abstract compiler to a micro-policy machine? and (2) can we support a stronger attacker model by protecting mutually distrustful

components against each other?

We are the first to work on question 1, and among the first to study question 2: Micro-policies are a recent hardware mechanism [9, 14] that is flexible and fine-grained enough to allow building a secure compiler against this strong attacker model. In independent parallel work [31,34], Patrignani et al. are trying to extend their previous results [32] to answer question 2 using different mechanisms (e.g. multiple protected modules and randomization). Related work is further discussed in §7.

1.3 Our Contribution

In this work we propose a new attacker model for secure compilation that extends the well-known notion of full abstraction to ensure protection for mutually distrustful components (§2). We devise a secure compilation solution (§4) for a simple object-oriented language (§5.1) that defends against this strong attacker model. Our solution includes a simple compiler chain (compiler, linker, and loader; §6.1) and a novel micro-policy (§6.2) that protects the abstractions of our simple language—class isolation, the method call discipline, and type safety—against arbitrary low-level attackers. Enforcing a method call discipline and type safety using a micro-policy is novel and constitutes a contribution of independent interest.

We have started proving that our compiler is secure, but since that proof is not yet finished, we do not present it in the report. Section 8.2 explains why we have good hopes in the efficiency and transparency of our solution for the protection of realistic programs. We also discuss ideas for mitigation when our mechanism is not transparent enough. However, in both cases gathering evidence through experiments to confirm our hopes is left for future work.

1.4 Other Insights

Throughout this work, we reasoned a lot about abstractions. One insight we gained is that even very simple high-level languages are much more abstract than one would naively expect. Moreover, we learned that some abstractions — such as functional purity — are impossible to efficiently enforce dynamically.

We also needed to extend the current hardware and formalism of micro-policies (§5.3) in order to achieve our challenging security goal. We needed two kinds of extensions: some only ease micro-policy writing, while the others increase the power to the monitoring mechanism. The first ones require a policy compiler, allowing an easier specification for complex micro-policies, which can then still run on the current hardware. The second ones require actual hardware extensions. Both of these extensions keep the spirit of micro-policies unchanged: Rules, in particular, are still specified as a mapping from tags to tags.

Finally, as we mention in §6.1.2, we were able to provide almost all security at the micro-policy level rather than the compiler level. This is very encouraging because it means that we might be able to provide full abstraction for complex compilers that already exist, using micro-policies while providing very little change to the compiler itself.

2 Stronger Attacker Model for Secure Compilation of Mutually Distrustful Components

Previous work on secure compilation [4,19,22,32] targets a property called full abstraction [1]. This section presents full abstraction (§2.1), motivates why it is not enough in the context of mutually distrustful components (§2.2), and introduces a stronger attacker model for this purpose (§2.3).

2.1 Full Abstraction

Full abstraction is a property of compilers that talks about the observable behaviors of partial programs evolving in a context. When we use full abstraction for security purposes, we will think of contexts as *attackers* trying to learn the partial program's secrets, or to break its internal invariants. Full abstraction relates the observational power of low-level contexts to that of high-level contexts. Hence, when a compiler achieves full abstraction, low-level attackers can be modeled as high-level ones, which makes reasoning about the security of programs much easier: Because they are built using source-level abstractions, high-level attackers have more structure and their interaction with the program is limited to that allowed by the semantics of the source language.

In order to state full abstraction formally, one first has to provide definitions for partial programs, contexts, and observable behaviors both in the high- and the low-level. Partial programs are similar to usual programs; but they could still be missing some elements—e.g. external libraries—before they can be executed. The usual, complete programs can be seen as a particular case of partial programs which have no missing elements, and are thus ready for execution.

A context is usually defined as a partial program with a hole; this hole can later be filled with a partial program in order to yield a new partial program. Finally, observable behaviors of complete programs can vary depending on the language and may include, termination, I/O during execution, final result value, or final memory state.

The chosen definition for contexts will set the granularity at which the attacker can operate. Similarly, defining the observable behaviors of complete programs can affect the observational power of the attacker in our formal model. The attacker we want to protect the program against is the context itself: The definition we choose for observable behaviors should allow the context to produce an observable action every time it has control, thus letting it convert its knowledge into observable behaviors. In our case, our source and target languages feature immediate program termination constructs. We can thus choose *program termination* as an observable behavior which models such strong observational power.

We denote high-level partial programs by P, Q, and high-level contexts by A. We denote by $A[P]$ the partial program obtained by inserting a high-level partial program P in a high-level context A. We denote low-level partial programs by p, q, and high-level contexts by a. We denote by $a[p]$ the insertion of a low-level partial program p in a low-level context a. Given a high-level partial program P, we denote by $P{\downarrow}$ the low-level program obtained by compiling P. We denote the fact that two complete high-level programs P and Q have the same observable behavior by $P \sim_H Q$. For two complete low-level programs p and q, we denote this by $p \sim_L q$. With these notations, full abstraction of the compiler is stated as

$$(\forall\, A, A[P] \sim_H A[Q]) \iff (\forall\, a, a[P{\downarrow}] \sim_L a[Q{\downarrow}])$$

for all P and Q. Put into words, a compiler is fully abstract when any two high-level partial programs P and Q behave the same in every high-level context if and only if the compiled partial programs $P{\downarrow}$ and $Q{\downarrow}$ behave the same in every low-level context. In other words, a compiler is fully abstract when a low-level attacker is able to distinguish between exactly the same programs as a high-level attacker.

Intuitively, in the definition of full abstraction the trusted compiled program ($P{\downarrow}$ or $Q{\downarrow}$) is protected from the untrusted low-level context (a) in a way that the context cannot cause more harm to the program than a high-level context (A) already could. This static separation between the trusted compiled program and the context is in practice chosen by the user and communicated to the compiler chain, which can insert a single protection barrier between between the two. In particular, in the definition of full abstraction the compiler is only ever invoked for the protected program ($P{\downarrow}$ or $Q{\downarrow}$), and can use this fact to its advantage, e.g. to add dynamic checks. Moreover, the definition of $a[p]$ (low-level linking) can insert a dynamic protection barrier between the trusted p and the untrusted a. For instance, Patrignani et al. [32] built a fully abstract compilation scheme targeting protected module architectures by putting the compiled program into a protected part of the memory (the protected module) and giving only unprotected memory to the context. A *single* dynamic protection barrier is actually enough to enforce the full abstraction attacker model.

2.2 Limitations of Full Abstraction

We study languages for which programs can be decomposed into *components*. Real-world languages have good candidates for such a notion of components: depending on the granularity we target, they could be packages, modules, or classes. Our compiler is such that source components are *separately compilable* program units, and compilation maps source-level components to target-level components.

When using a fully abstract compiler in the presence of multiple components, the user has a choice whether a component written in the high-level language is trusted, in which case it is considered part of the program, or untrusted, in which case it is considered part of the context. If it is untrusted, the component can as well be compiled with an insecure compiler, since anyway the fully abstract compiler only provides security to components that are on the good side of the protection barrier. If the program includes components written in the low-level language, e.g. for efficiency reasons, then the user has generally no choice but to consider these components untrusted. Because of the way full abstraction is stated, low-level components that are not the result of the compiler cannot be part of the trusted high-level program, unless they have at least a high-level equivalent (we discuss this idea in §2.3).

Figure 1 graphically illustrates how full abstraction could be applied in a multi-component setting. Components **C1**, **C2**, and **C3** are written in the high-level language, while **c4** and **c5** are written in the low-level one. Suppose the user chooses to trust **C1** and **C2** and not to trust **C3**, then the compiler will introduce a single barrier protecting **C1** and **C2** from all the

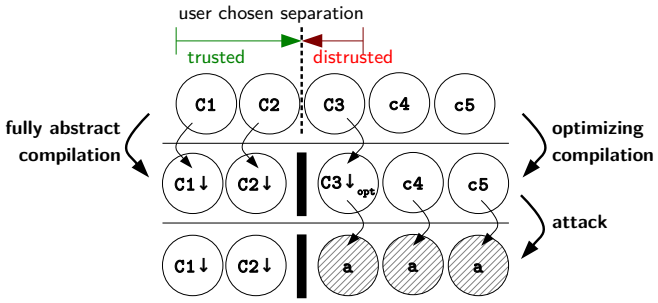

Figure 1: Full abstraction for multiple components

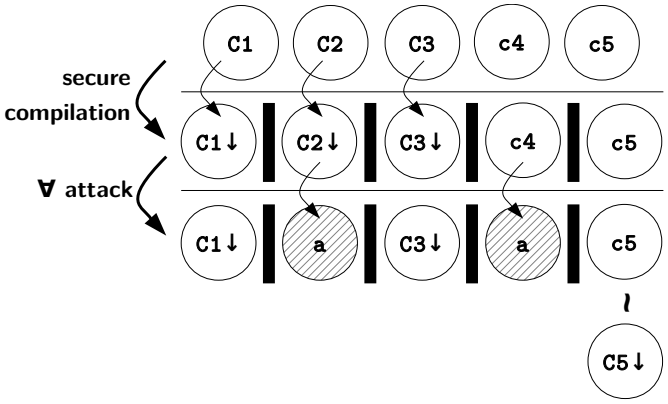

Figure 2: Secure compilation for mutually distrustful components

other components.

There are two assumptions on the attacker model when we take full abstraction as a characterization of secure compilation: *the user correctly identifies trusted and untrusted components so that (1) trusted components need not be protected from each other, and (2) untrusted components need no protection whatsoever.* We argue that there are common cases in which the absolute, binary trust notion implied by full abstraction is too limiting (e.g. there is no way to achieve all the user's security goals), and for which a stronger attacker model protecting mutually distrustful components is needed.

Assumption (1) is only realistic if all trusted components are memory safe [13] and do not exhibit C-style undefined behaviors. Only when all trusted components have a well-defined semantics in the high-level language is a fully abstract compiler required to preserve this semantics at the low level. Memory safety for the trusted components may follow either from the fact that the high-level language is memory safe as a whole or that the components have been verified to be memory safe [5]. In the typical case of unverified C code, however, assumption (1) can be unrealistically strong, and the user cannot be realistically expected to decide which components are memory safe. If he makes the wrong choice all bets are off for security, a fully abstract compiler can produce code in which a control hijacking attack [15, 41] in one trusted component can take over all the rest. While we are not aware of any fully abstract compiler for unverified C, we argue that if one focuses solely on achieving the full abstraction property, such a compiler could potentially be as insecure in practice as standard compilers.

Even in cases where assumption (1) is acceptable, assumption (2) is still a very strong one. In particular, since components written in the low-level language cannot get protection, every security-critical component would have to be written in the high-level source language, which is often not realistic. Compiler correctness would be sufficient on its own if all components could be written in a safe high-level language. The point in moving from compiler correctness to full abstraction, which is stronger, is precisely to account for the fact that some components have to be written in the low-level language, e.g. for performance reasons.

Assumption (2) breaks as soon as we consider that it makes a difference whether the attacker owns one or all the untrusted components. As an example, assume that an attacker succeeds in taking over an untrusted component that was used by the program to render the picture of a cat. Would one care whether this allows the attacker to also take over the low-level cryptographic library that manages private keys? We believe that the cryptographic library, which is a security-critical component, should get the same level of protection as a compiled component, even if for efficiency it is implemented in the low-level language.

When assumption (1) breaks, trusted components need to be protected from each other, or at least from the potentially memory unsafe ones. When assumption (2) breaks, untrusted security-critical components need to be protected from the other untrusted components. In this work, we propose a stronger attacker model that removes both these assumptions by requiring all components to be protected from each other.

2.3 Mutual Distrust Attacker Model

We propose a new attacker model that overcomes the previously highlighted limitations of full abstraction. In this attacker model, we assume that each component could be compromised and protect all the other components from it: we call it an attacker model for

mutually distrustful components. This model can provide security even in C-like unsafe languages when some of the high-level components are memory unsafe or have undefined behaviors. This is possible if the high-level semantics treats undefined behavior as arbitrary changes in the state of the component that triggered it, rather than in the global state of the program. In the following we will assume the high-level language is secure.

All compiled high-level components get security unconditionally: the secure compiler and the dynamic barriers protect them from all low-level attacks, which allows reasoning about their security in the high-level language. For low-level components to get security they have to satisfy additional conditions, since the protection barriers are often not enough on their own for security, as the compiler might be inserting boundary checks, cleaning registers, etc. and the low-level code still needs to do these things on its own in order to get full protection. Slightly more formally, in order for a low-level component **c** to get security it must behave in all low-level contexts like some compiled high-level component **C**↓. In this case we can reason about its security at the high level by modelling it as **C**. This captures the scenario in which **c** is written in the low-level language for efficiency reasons.

We illustrate our stronger attacker model in figure 2. The protected program is the same as in the previous full abstraction diagram of figure 1. This time, however, the user doesn't choose a trust barrier: all components are considered mutually distrustful instead. Each of them gets protected from the others thanks to barriers inserted by the compiler. While components **C3**, **c4**, and **c5** were distrusted and thus not protected in the previous diagram, here all of them can get the same amount of protection as other components. To get security **C3** is compiled using the secure compiler, while for **c4** and **c5** security is conditioned on equivalence to high-level components; in the figure we assume this only for **c5**. The attacker can compromise arbitrary components (including high-level compiled components), e.g. **C2**↓ and **c4** in the diagram. In this compromise scenario, we ensure that the uncompromised components **C1**↓, **C3**↓, and **c5** are protected from all low-level attacks coming from the compromised components. In general, our attacker model defends against all such compromise scenarios.

To sum up, our attacker model can be stated as follows: (a) the attacker compromises with component granularity, (b) the attacker may compromise any set of components, (c) in every compromise scenario, each uncompromised compiled high-level component is secure against low-level attacks from all compromised components, and (d) in every compromise scenario, each uncompromised low-level component that has a high-level equivalent is secure against low-level attacks from all compromised components.

3 Micro-Policies and the PUMP: Efficient Tag-Based Security Monitors

We present *micro-policies* [9, 14], the mechanism we use to monitor low-level code so as to enforce that our compiler is secure. Micro-policies [9, 14] are a tag-based dynamic protection mechanism for machine code. The reference implementation on which micro-policies are based is called the PUMP [14] (Programmable Unit for Metadata Processing).

The PUMP architecture associates each piece of data in the system with a *metadata tag* describing its provenance or purpose (e.g. "this is an instruction," "this came from the network," "this is secret," "this is sealed with key k"), propagates this metadata as instructions are executed, and checks that policy rules are obeyed throughout the computation. It provides great flexibility for defining policies and puts no arbitrary limitations on the size of the metadata or the number of policies supported. Hardware simulations show [14] that an Alpha processor extended with PUMP hardware achieves performance comparable to dedicated hardware on a standard benchmark suite when enforcing either memory safety, control-flow integrity, taint tracking, or code and data separation. When enforcing these four policies simultaneously, monitoring imposes modest impact on runtime (typically under 10%) and power ceiling (less than 10%), in return for some increase in energy usage (typically under 60%) and chip area (110%).

The reference paper on micro-policies [9] generalizes previously used methodology [8] to provide a generic framework for formalizing and verifying *arbitrary* policies enforceable by the PUMP architecture. In particular, it defines a *generic symbolic machine*, which abstracts away from low-level hardware details and serves as an intermediate step in correctness proofs. This machine is parameterized by a *symbolic micro-policy*, provided by the micro-policy designer, that expresses tag propagation and checking in terms of structured mathematical objects rather than bit-level concrete representations. The micro-

policies paper also defines a *concrete machine* which is a model of PUMP-like hardware, this time including implementation details.

The proposed approach to micro-policy design and verification is presented as follows. First, one designs a reference *abstract machine*, which will serve as a micro-policy specification. Then, one instantiates the generic symbolic machine with a symbolic micro-policy and proves that the resulting symbolic machine *refines* the abstract machine: the observable behaviors of the symbolic machine are also legal behaviors of the abstract machine, and in particular the symbolic machine fail-stops whenever the abstract machine does. Finally, the symbolic micro-policy is implemented in low-level terms, and one proves that the concrete machine running the micro-policy implementation refines the symbolic machine.

In this work, we use a slightly modified symbolic machine as the target of our secure compiler. Our symbolic machine differs from the previous one [9] in two ways: First, its memory is separated into regions, which are addressed by *symbolic pointers*. Note, however, that our protection does not rely on this separation but only on the tagging mechanism itself: in particular, all components can refer to all symbolic pointers, without restriction. Mapping memory regions to concrete memory locations before executing the program on a concrete machine would be a main task of a *loader* — we leave a complete formalization of loading for future work. Second, we extend the micro-policy mechanism itself, allowing rules to read and write the tags on more components of the machine state. We detail these extensions in section 5.3, which is dedicated to our target machine. We also leave for future work the implementation of these additional tags in the PUMP rules, their formalization in an extended concrete machine, and the extension of our results to the concrete level.

4 Compilation Chain Overview

In this section, we present our compiler and give intuition about the connections between the different parts that play a role in our solution.

Our compilation chain, which we present in figure 3, splits into a two-step compiler, a linker and a loader. It produces a program to execute on the symbolic micro-policy machine. Our dedicated protection micro-policy will be loaded into the machine, allowing proper runtime monitoring of the program.

The compilation chain takes source components (e.g. a main program and standard libraries) and target components (e.g. low-level libraries) as input, and outputs a target-executable program. Components must all come with interface specifications, written in a common interface language. These interfaces specify the definitions that each component provides, and the definitions it expects other components to provide.

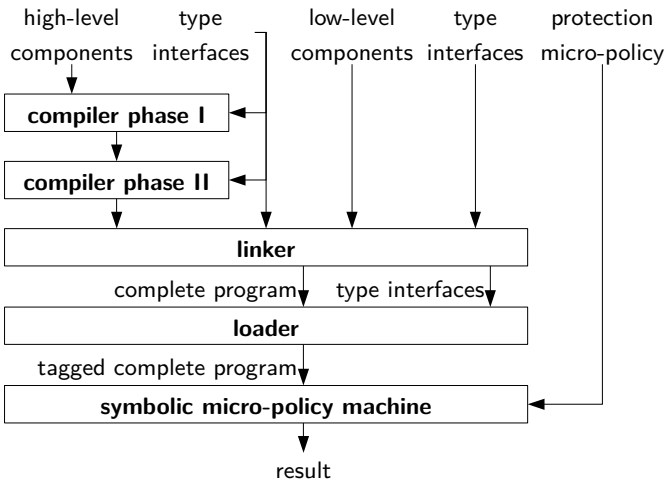

Figure 3: Overview of the compilation chain

In the compilation phase, the compiler first translates source components to an intermediate representation, which then gets translated to target components.

In the linking phase, the linker checks that the interfaces of the components getting linked are compatible. It also makes sure that all components only give definitions under names that follow from their interfaces; and symmetrically that they do provide a definition for each of these names. Is so, the linker puts them together to form a partial program, and makes sure that this partial program is actually complete (i.e. no definition is missing).

In the loading phase, the loader builds an initial machine state out of this complete program by tagging its memory using type information that was gathered by the linker. The result is thus a target-level executable program — i.e. a symbolic machine tagged program. The loader's tagging will force all components to correctly implement the interfaces that was provided for them, when we later run and monitor them using our protection micro-policy: The machine will failstop as soon as a component violates its interface upon interaction with another component (violations on internal computational steps is not harmful and hence allowed).

Because we required that every component should have an interface, low-level libraries that were com-

piled using other compilers — e.g. C compilers — are only usable in our system if somebody writes interfaces for them. In section 8 we discuss more generally the need for manual or automated *wrapping* of low-level code; once we have a way to wrap low-level code, providing a matching interface will be easy.

5 Languages and Machines

In this section, we present and formalize the languages that are used in the compilation scheme. We introduce a simple object-oriented language, an abstract stack machine, and an extended symbolic micro-policy machine with segmented memory. The first will be our source language and the last our target machine, while the intermediate abstract machine will offer a different view from the source language which makes the connection with the low level more direct. The source language includes constructs for specifying the interfaces of components; these get reused as-is at all levels.

5.1 Source Level: An Object-Oriented Language

We first formalize our source language, beginning with the interface constructs that it offers and then presenting its syntax and semantics. The source language we consider is an imperative class-based object-oriented language with static objects, private fields, and public methods. It is inspired by previous formalizations of Java core calculi [10, 21] and Java subsets [23], with the aim of keeping things as simple as possible. As a consequence, we do not handle inheritance nor dynamic allocation, which are available in all these works.

We start with the simplicity of Featherweight Java [21], and add imperative features in the way Middleweight Java [10] does. However, we do not add as many imperative features: just branching, field update and early termination (the latter is not a feature of Middleweight Java). The resulting language is similar to Java Jr. [23] with early termination, but without packages: Our granularity for components is that of classes instead.

Example components which encode some usual types are provided in appendix section A.1, and could help in getting familiar with this language.

IDT ::= — import declaration tables
DT

EDT ::= — export declaration tables
DT

DT ::= — declaration tables
(CDT, ODT)

CDT ::= — class declaration table
$[] \mid \{c \mapsto CD\} :: CDT$

CD ::= — class declaration
$\mathbf{class\,decl}\,\{MD_1, \ldots, MD_k\}$

MD ::= — method declaration
$c_{\mathsf{r}}(c_{\mathsf{a}})$

ODT ::= — object declaration table
$[] \mid \{o \mapsto OD\} :: ODT$

OD ::= — object declaration
$\mathbf{obj\,decl}\,c$

Figure 4: Interface language syntax

5.1.1 Interfacing: A Specification Language for Communicating Classes

The notion of component in this work is that of a class c together with all its object instances' definitions. Because we have no dynamic allocation, for the moment these instances are simply all the static objects defined with type c. To allow interaction between components while being able to separately compile them, we have a simple interface syntax based on import and export declarations. This interface language gives the external view on source components and is presented in figure 4.

Syntax and Naming Conventions Object names o and class names c are global and assumed to be arbitrary natural numbers. The two main syntactic constructs in the interface language are class declarations and static object declarations. Class declarations specify public methods with argument and result types without providing their body; no fields are declared in an interface because we only consider private fields in our languages. Static object declarations specify an object to be an instance of a given class, without providing the values of its fields.

The interface of a partial program at all levels is composed of an import declaration table IDT speci-

fying the class and object definitions it expects other components to provide, and an export declaration table EDT which declares the definitions that this partial program provides. Export and import declaration tables share common syntax and are composed of class and object declarations. The type of the declared objects must come from the classes specified by the partial program: defining object instances of classes coming from other components is not allowed. Intuitively, our object constructors (and fields) are private and the only way to interact with objects is via method calls.

In contrast with objects and classes to which we refer using global names, methods are class-wise indexed: the methods m of a class c are referred to as $1, \ldots, k$ following the order of their declarations. (The same goes for fields f, below.) The syntax we consider for names can be thought of as coming out of a parser, that would take a user-friendly Java-like syntax and perform simple transformations so that the names match our conventions.

Use at Linking and Loading We have presented the roles of the linker and the loader when we introduced the compilation chain in section 4: We can model linking as an operation that takes two target partial programs and their interfaces, and yields a new partial program which contains all definitions from both partial programs, with a new matching interface. Loading then takes a complete target program and tags it, yielding a machine state that can be reduced using the semantics of our symbolic micro-policy machine. Let us now explain how interfaces are used at linking and loading.

A class import declaration gives the precise type signatures that the partial program expects from the methods of the corresponding class: When linking against a partial program that defines this class, the class export declaration should exactly match with the import one. Similarly, an import object declaration gives the expected type for this object, and should match with the corresponding export declaration when linking against a partial program that defines it.

Two partial programs have compatible interfaces when (1) they don't both have export declarations for the same class nor the same object, and (2) for every name in an import declaration of any of the two programs, if the other program happens to have an export declaration for this name, then the import and export declarations are syntactically equal. Linking two partial programs with compatible interfaces yields a new partial program with updated import/export declarations: Export declarations are combined, and import declarations that found matching export declarations are removed. When all partial programs have been linked together, the linker can check that there are no remaining import declarations to make sure that the program is complete.

Finally, the loader will make use of the export declarations to ensure that programs comply with the export declarations they declared: In the untyped target language, the loader sets the initial memory tags in accordance with the export declarations, which will allow our micro-policy to perform dynamic type checking. This will be further explained in section 6.2.3.

$$
\begin{array}{ll}
P,\ Q ::= & \text{source program} \\
\quad (IDT, T, EDT) & \\
T ::= & \text{definition tables} \\
\quad (CT, OT) & \\
CT ::= & \text{class definition table} \\
\quad [] \mid \{c \mapsto C\} :: CT & \\
C ::= & \text{class definition} \\
\quad \mathbf{class}\, \{c_1, \ldots, c_k; M_1, \ldots, M_l\} & \\
M ::= & \text{method definition} \\
\quad c_{\mathsf{r}}(c_{\mathsf{a}})\{e\} & \\
e ::= & \text{expressions} \\
\quad \mathbf{this} \mid \mathbf{arg} \mid o \mid e.f \mid e.f := e' \mid e.m(e') & \\
\mid e == e' \ ?\ e'' : e''' \mid \mathbf{exit}\ e \mid e; e' & \\
OT ::= & \text{object definition table} \\
\quad [] \mid \{o \mapsto O\} :: OT & \\
O ::= & \\
\quad \mathbf{obj}\, c\{o_1, \ldots, o_k\} &
\end{array}
$$

Figure 5: Source language syntax

5.1.2 Source Syntax and Semantics

The syntax of our source language is presented in figure 5. The two main syntactic constructs in this language are class definitions and static object definitions. Class definitions declare typed private fields and define public methods with argument and result types as well as an expression which serves as a method body. Static object definitions define instances of defined classes by providing well-typed val-

ues for the fields. For simplicity, methods only take one argument: this does not affect expressiveness because our language is expressive enough to encode tuple types (appendix section A.1 shows examples that encode more complex types than tuple types).

Most expressions are not surprising for an object-oriented language: apart from static object references o and variables (**this** for the current object and **arg** for the method's argument), we have support for selecting private fields, calling public methods, and testing object identities for equality. The language also features field update and early termination. Both are crucial for modeling realistic low-level attackers in our high-level language: Low-level attackers can indeed keep information between calls using the memory and stop the machine whenever they have control. We thus add primitives that can model this in the high-level: field updates enable methods to have state (methods are not pure functions anymore), and early termination allows an attacker to prematurely end the program.

Like we already mentioned, fields are private and methods are public. This means that in the method body of a given class, the only objects whose fields may be accessed are the instances of that specific class. The only way to interact with object instances of other classes is to perform a method call.

The only values in the language are object names, and the only types are class names. The language comes with a type system that ensures that object and method definitions match with the types that were declared for them. Our language does not feature dynamic allocation, inheritance, or exceptions. We hope to add some of these features in the future. Loops are simulated using recursive method calls and branching is done via object identity tests.

The semantics of the source language is standard and is given in appendix A.2.

5.2 Intermediate Level: An Object-Oriented Stack Machine

Our intermediate machine is an object-oriented stack machine with one local stack per class. The syntax for intermediate machine programs is presented in figure 6. The main syntactic construct is that of a compartment, which is the notion of component at this level. A compartment combines a class definition with all the object instances of this class and with a local stack.

The main difference with respect to the source language is that instead of expressions, method bodies

$IP,\ IQ ::=$ intermediate program
(IDT, ICT, EDT)

$ICT ::=$ compartment table
$[]\ \mid \{c \mapsto IC\} :: ICT$

$IC ::=$ compartment declaration
$\{IM_1, \ldots, IM_k; LOT; LS\}$

$IM ::=$ method body
$Icode$

$Icode ::=$ intermediate machine code
$[]\ \mid Iinstr; Icode$

$Iinstr ::=$ machine instruction
$\mathbf{Nop} \mid \mathbf{This} \mid \mathbf{Arg} \mid \mathbf{Ref}\, o \mid \mathbf{Sel}\, f \mid \mathbf{Upd}\, f$
$\mid \mathbf{Call}\, c\, m \mid \mathbf{Ret} \mid \mathbf{Skip}\, n \mid \mathbf{Skeq}\, n \mid \mathbf{Drop} \mid \mathbf{Halt}$

$LOT ::=$ compartment local object table
$[]\ \mid \{o \mapsto LO\} :: LOT$

$LO ::=$ local instance definition
$(o_1, \ldots, o_l)$

$LS ::=$ local stack
$[]\ \mid o :: LS$

Figure 6: Intermediate language syntax

are expressed as sequences of instructions in abstract machine code. These instructions manipulate values stored on the local stack associated with each compartment.

Nop does nothing. **This**, **Arg** and **Ref** o put an object on the local stack — the current object for **This**, the current method argument for **Arg**, and object o for **Ref** o. **Sel** f pops an object from the stack, selects field f of this object and pushes the selected value back to the stack. **Upd** f pops a value and an object, sets the f field of the object to the popped value, then pushes back this value on the stack. **Call** $c\ m$ pops an argument value and a target object o, performs a call of the object o's m method with the popped argument if o has type c (otherwise, the machine failstops). The callee can then use the **Ret** instruction to give control back to the caller: this instruction pops a value from the callee's stack and pushes it on the caller's stack, as a result value for the call. **Skip** n skips the n next instructions. **Skeq** n pops two values from the local stack and either skips the n next instructions if the values are equal, or does nothing more if they are different. **Drop** removes the top value from the stack. **Halt** halts the machine immediately, the result of the computation being the current top value on the stack.

The purpose of this intermediate language is to ease going low-level. In particular, its syntax with the **Call** instruction being annotated with a class makes explicit the fact that method calls are statically resolved by the source to intermediate compiler. This is possible in our source language, because we have no inheritance.

mem ::= memory
$[] \mid \{loc \mapsto R\} :: mem$

loc ::= region symbolic location
$\mathbf{objl}\ o \mid \mathbf{methl}\ c\ m \mid \mathbf{stackl}\ c$

R ::= tagged memory region
$[] \mid (word\ @\ t_{\mathsf{mem}}) :: R$

$word$ ::= symbolic machine word
$n \mid loc + n \mid \mathbf{encode}\ instr$

$instr$ ::= machine instruction
$\mathbf{Nop} \mid \mathbf{Const}\ i\ r_{\mathsf{d}} \mid \mathbf{Mov}\ r_{\mathsf{s}}\ r_{\mathsf{d}} \mid \mathbf{Binop}_{op}\ r_1\ r_2\ r_{\mathsf{d}} \mid \mathbf{Load}\ r_{\mathsf{p}}\ r_{\mathsf{d}} \mid \mathbf{Store}\ r_{\mathsf{p}}\ r_{\mathsf{s}} \mid \mathbf{Jump}\ r \mid \mathbf{Jal}\ r \mid \mathbf{Bnz}\ r\ i \mid \mathbf{Halt}$

i ::= immediate value
$n \mid loc + n$

Figure 7: Symbolic machine memory

5.3 Target Level: An Extended Symbolic Micro-Policy Machine

Here, we present our the target of our compiler: an extended symbolic micro-policy machine with segmented memory. We first recall the common basis for our machine and the symbolic machine presented in the original micro-policies paper [9], then present and comment on the differences between them.

5.3.1 Symbolic Micro-Policy Machine

A symbolic micro-policy machine [9] is an executable model of micro-policies that abstracts away from some of the implementation details (e.g. the implementation of the micro-policy monitor in machine code). The definition of a symbolic micro-policy machine is abstracted over a *symbolic micro-policy*.

$LP,\ LQ$ ::= low-level program
$(IDT, LPmem, EDT)$

$LPmem$ ::= program memory
$[] \mid \{loc \mapsto LPR\} :: LPmem$

LPR ::= program memory region
$[] \mid word :: LPR$

i ::= immediate value
$n \mid loc + n$

Figure 8: Symbolic machine program syntax

In our case, a symbolic micro-policy is defined as a collection of *symbolic tags*, which are used to label instructions and data, and a *transfer function*, which is invoked on each step of the machine to propagate tags from the current machine state to the next one. We ignore *monitor services* of Azevedo de Amorim et al. [9] and *extra pieces of state* which are only used by monitor services — because we don't need them: we successfully avoid monitor services, which in the context of micro-policies are much more expensive than rules. The transfer function is a mathematical function that defines the micro-policy rules — in the mechanized metatheory of the original micro-policies paper this function is written in Gallina, the purely-functional language at the heart of the Coq proof assistant.

A machine state of our symbolic micro-policy machine is composed of a memory, a register file of general-purpose registers, and a program counter register *pc*. The program counter register points to a location in memory which contains the next instruction to execute.

We present the list of instructions in figure 7, together with other memory-related definitions on which we will focus when we later explain the segmented memory. These instructions are those of the machine code of the original micro-policies paper [9]: **Nop** does nothing. **Const** $i\, r_{\mathsf{d}}$ puts an immediate constant i into register r_{d}. **Mov** $r_{\mathsf{s}}\, r_{\mathsf{d}}$ copies the contents of r_{s} into r_{d}. **Binop**$_{op}\, r_1 r_2 r_{\mathsf{d}}$ performs a binary operation *op* (e.g. addition, subtraction, equality test) on the content of registers r_1 and r_2, and puts the result in register r_{d}. **Load** $r_{\mathsf{p}}\, r_{\mathsf{s}}$ copies the content of the memory cell at the memory location stored in r_{p} to r_{s}. **Store** $r_{\mathsf{p}}\, r_{\mathsf{s}}$ copies the content of r_{s} to the memory cell at the memory location stored in r_{p}. **Jump** and **Jal** (jump-and-link) are unconditional indirect jumps, while **Bnz** $r\, i$ branches to a fixed offset *imm* (relative to the current pc) if register r is nonzero. **Halt** halts the machine immediately.

In the following, we denote **Binop**$_+ r_1 r_2 r_{\mathsf{d}}$ (addition) by **Add** $r_1\, r_2\, r_{\mathsf{d}}$, and **Binop**$_- r_1 r_2 r_{\mathsf{d}}$ (subtraction) by **Sub** $r_1\, r_2\, r_{\mathsf{d}}$.

5.3.2 Extensions to Monitoring Mechanism

We consider a more powerful symbolic micro-policy machine that allows the transfer function to inspect and update more tags.

First, we assume that the transfer function produces new tags for the input arguments of the instructions; not only for the output one. This is required, for instance, to transfer a linear capability from an input to an output: one wants not only to copy the capability in the output tag, but also to erase it from the input tag.

Second, we assume that there are some fixed registers whose tags can always be checked and updated by the transfer function, even if the registers are neither input nor output to the current instruction. This allows us to efficiently clean these fixed registers upon calls and returns.

Third, we assume that the transfer function receives as an argument not only the tag of the current instruction, but also the tag on the *next* instruction. For instance, when monitoring a **Jump** instruction, we assume the tag on the targeted location can be checked. This extension allows us to write and explain our micro-policy in a much simpler way.

The first two extensions require extra hardware support. For the last extension, however, we conjecture that our micro-policy — and probably other similar micro-policies — can be transformed into a policy which doesn't have to perform next instruction checks. This kind of translation has already been done by hand, for example in a previous compartmentalization micro-policy [9]: the next instruction checks are replaced by current instruction checks happening on the next step, making the machine failstop one step later in the case of an error. We leave for future work the writing of a policy compiler doing such a transformation automatically.

5.3.3 Easing Component-Oriented Reasoning: Segmented Memory

Instead of having a monolithic word-addressed memory, the machine we consider has a segmented memory which consists of several memory regions which are addressed by *symbolic locations*. Targeting such a machine allows for easy separate compilation, and is a pragmatic intermediate step when going towards real machine code, which we plan to do in the future.

As presented in figure 7, the definition of memory directly mentions symbolic locations. A generic symbolic machine definition would be abstracted over the definition of symbolic locations, but in our case, we define them to be either method, object, or stack locations for reasons that will be clear when we present our compiler in section 6.1.2. Our instantiation of memory tags t_{mem} will be studied with other definitions related to the symbolic micro-policy, in section 6.2.

Immediate constants and words in the symbolic

machine are extended to include symbolic locations with an offset, which are memory addresses: The k memory cells of a region located at *loc* are addressed from $loc + 0$ to $loc + (k - 1)$. In the following, we use the simpler notation *loc* for $loc + 0$.

Words are also extended to include a new **encode** *instr* construct: Decoding instructions in this machine is a symbolic operation of deconstructing such a word. Now that instructions feature symbolic locations with an offset as immediate values, it would indeed have no practical meaning to try to extend the encoding of instructions to this. When we use the PUMP in future work, some of the symbolic-level instructions could have to be compiled to sequences of PUMP instructions: PUMP memory addresses, the PUMP equivalent of symbolic locations with an offset, are word-sized and thus too big to fit in a PUMP immediate value. Another solution would be to restrict addressable memory to less than a full word; the symbolic encoding allows us to delay the choice between these solutions.

Tags are not affected by all the changes that we listed, hence the monitoring mechanism isn't affected either. The semantics, however, is affected in the following way: Trying to decode an instruction out of a regular word or of a symbolic location with an offset failstops the machine. All instructions failstop when one of their operands has the form **encode** *instr*. A **Jump**, **Jal**, **Load** or **Store** instruction used with a regular word for the pointer failstops the machine. These instructions keep their usual behavior when provided with a symbolic location and a valid offset; if it does not correspond to a valid memory address, however, the machine failstops. Most binary operations failstop when the first or second operand is a symbolic location with an offset: exceptions are 1) addition and subtraction with regular words, when the first operand is the location with an offset, which respectively increment and decrement the offset accordingly; 2) equality tests between symbolic locations with offsets. Finally, the **Bnz** instruction failstops when the provided register or immediate value is a symbolic location with an offset.

The syntax for symbolic machine programs is presented in figure 8. They define a memory which is like the one of the symbolic machine, except that cells are not tagged: The tags for the memory will be provided by the low-level loader, which will be detailed in the next section.

6 Our Solution: Protecting Compiled Code with a Micro-Policy

In this section, we present our solution for the secure compilation of mutually distrustful components: first we describe our simple two-step compiler, then we present our micro-policy dynamically protecting components from each other.

6.1 Two-Step Compilation

We start with our two-step compilation: first the compilation of source programs to intermediate machine programs, then the one of intermediate machine programs to target machine programs.

6.1.1 From Source to Intermediate

Our type-preserving source to intermediate compiler is a mostly direct translation of source expressions to abstract machine instructions, which gives a lower-level view on source components. Nothing in the translation is related to security: we provide security at this level by giving appropriate semantics to intermediate-level instructions, which make them manipulate local stacks and local object tables rather than a single global stack and a single object table. In this translation, we statically resolves method calls, which is possible because our language doesn't feature inheritance.

A high-level component is easily mapped to an intermediate compartment: Method bodies are compiled one by one to become intermediate-level method bodies. Object definitions corresponding to that component are taken from the global object table OT and put in the local object table LOT of the compartment. Finally, the stack is initialized to an empty stack.

Compiling Source Expressions to Stack Machine Code Assuming that method m of class c in program P has definition $c_{\mathsf{r}}(c_{\mathsf{a}})\{e\}$, compilation is defined as follows:

$$\mathscr{C}(P, c, m) = \mathscr{A}(e); \mathbf{Ret}$$

The $\mathscr{A}$ function is recursively defined as presented in figure 9; we allow ourselves to refer to P, c and m in this definition. We denote by $P; c, m \vdash e : c'$ the predicate indicating that expression e has type c' when typed within method m of class c from program P. In the compilation of method calls, we assume a type inference algorithm which, given P, c and m,

$$
\begin{aligned}
\mathscr{A}(\mathbf{this}) &= \mathbf{This} \\
\mathscr{A}(\mathbf{arg}) &= \mathbf{Arg} \\
\mathscr{A}(o) &= \mathbf{Ref}\ o \\
\mathscr{A}(e.f) &= \mathscr{A}(e); \mathbf{Sel}\, f \\
\mathscr{A}(e.f := e') &= \mathscr{A}(e); \mathscr{A}(e'); \mathbf{Upd}\, f \\
\mathscr{A}(e.m'(e')) &= \mathscr{A}(e); \mathscr{A}(e'); \mathbf{Call}\, c'\, m' \\
&\quad \text{where } c' \text{ satisfying} \\
&\quad\ P; c, m \vdash e : c' \\
&\quad \text{is found by type inference} \\
\mathscr{A}(e_1 == e_2\ ?\ e_3 : e_4) &= \mathscr{A}(e_1); \mathscr{A}(e_2); \\
&\quad \mathbf{Skeq}\,(|\mathscr{A}(e_4)| + 1); \\
&\quad \mathscr{A}(e_4); \mathbf{Skip}\,|\mathscr{A}(e_3)|; \\
&\quad \mathscr{A}(e_3); \\
&\quad \mathbf{Nop} \\
\mathscr{A}(e; e') &= \mathscr{A}(e); \mathbf{Drop}; \mathscr{A}(e') \\
\mathscr{A}(\mathbf{exit}\ e) &= \mathscr{A}(e); \mathbf{Halt}
\end{aligned}
$$

Figure 9: Compiling source expressions to intermediate machine code

finds the unique type c' such that $P; c, m \vdash e : c'$. We use it to statically resolve method calls by annotating intermediate-level call instructions. In this document, we do not present the type system nor the type inference algorithm for our source language, which are standard.

The invariant used by the compilation is that executing $\mathscr{A}(e)$ in the intermediate level will either diverge—when evaluating e in the high-level diverges— or terminate with exactly one extra object on the local stack—in which case this object is exactly the result of evaluating e. In a method body, this object on top of the stack can then be returned as the result of the method call, which is why $\mathscr{A}(e)$ is followed by a **Ret** instruction in the main compilation function $\mathscr{C}$.

With this invariant in mind, the translation is rather straightforward, which is not surprising since our abstract stack machine was designed with this goal in mind. An important point is that we keep the evaluation order of the source language: left to right. It matters because of side effects and early termination being available in our language.

Let us explain the only non-trivial expression to compile: the object identity test ($e_1 == e_2\ ?\ e_3 : e_4$), for which we use two branching instructions $\mathbf{Skeq}\,(|\mathscr{A}(e_4)| + 1)$ and $\mathbf{Skip}\,|\mathscr{A}(e_3)|$. Here, we denote by $|\mathscr{A}(e)|$ the length of the sequence of instructions $\mathscr{A}(e)$. With equal objects, executing $\mathbf{Skeq}\,(|\mathscr{A}(e_4)|+1)$ will skip the code corresponding to e_4 and the $\mathbf{Skip}\,|\mathscr{A}(e_3)|$ instruction, hence branching directly to the code corresponding to e_3 to execute it. With different objects, executing $\mathbf{Skeq}\,(|\mathscr{A}(e_4)| + 1)$ will do nothing and execution will proceed with the code corresponding to e_4, followed by a $\mathbf{Skip}\,|\mathscr{A}(e_3)|$ instruction which will unconditionally skip the code corresponding to e_3, hence branching to the **Nop** instruction. The effect of all this is that in the end, we have executed e_1 and e_2 in this order, popped the resulting objects from the stack, and either executed e_3 if they were equal or e_4 if they were different: We execute the appropriate code in both cases.

6.1.2 From Intermediate to Target

We now present our unoptimizing, type-preserving translation from intermediate-machine compartments to target-level components. Target-level compartments are defined as sets of untagged symbolic machine memory regions. Like in the source to intermediate compilation, the translation itself is rather standard. The exception is that components cannot trust other components to follow conventions such as not messing with a global call stack, or not modifying some registers. As a consequence, components use local stacks and all relevant registers need to be (re)initialized when the compiled component takes control. Other than that, all security is enforced by means of instruction-level monitoring (§6.2).

Object Compilation Each object o that was assigned a definition $(o_1, \ldots, o_l)$ now gets its own region in target memory. This region is assigned symbolic location $\mathbf{objl}\, o$ and spans over l memory cells. These cells are filled with the $\mathbf{objl}\, o_1, \ldots, \mathbf{objl}\, o_l$ symbolic locations — which are the addresses of these objects in memory.

Local Stack Compilation Each local stack also gets its own memory region, under symbolic location $\mathbf{stackl}\, c$ where c is the name of the compartment being compiled. Components will maintain the following invariant during computation: The first cell holds a symbolic pointer to the top of the stack, which is $(\mathbf{stackl}\, c) + l$ where l is the length of the stack. The following cells are used for storing actual values on the stack.

Here, we only care about compiling intermediate-level components that come from the source to inter-

mediate compiler. For these components, the initial stack is always empty. Hence, we just initialize the first cell to **stackl** c, and the initial content of extra cells can be arbitrary constants: their only purpose is to be available for filling during the computation.

Method Compilation: From Coarse- to Fine-Grained Instructions A method with index m gets its own memory region under symbolic location **methl** $c\,m$ where c is the name of the compartment being compiled. The length of these memory regions is that of the corresponding compiled code, which is what they are filled with. The compilation is a translation of the method bodies, mapping each intermediate-level instruction to a low-level instruction sequence.

The compilation uses ten distinct general-purpose registers. Register r_{a} is automatically filled upon low-level call instructions **Jal** — following the semantics of the machine studied in the original micro-policies paper [9] — for the callees to get the address to which they should return. Registers r_{tgt}, r_{arg} and r_{ret} are used for value communication: r_{tgt} stores the object on which we're calling a method—the *target object*—and r_{arg} the argument for the method, while r_{ret} stores the result of a call on a return. Registers r_{aux1}, r_{aux2}, r_{aux3} are used for storing temporary results. Register r_{sp} holds a pointer to the current top value of the local stack — we call this register the *stack pointer* register. Register r_{spp} always holds a pointer to a fixed location where the stack pointer can be stored and restored – this location is the first cell in the memory region dedicated to the local stack. Finally, register r_{one} always stores the value 1 so that this particular value is always easily available.

The compilation of method m of class c with method body *Icode* is defined as follows:

$$
\begin{aligned}
&\mathscr{C}(c, m, \mathit{Icode}) = \\
&\quad \mathbf{Const}\ 1\ r_{\mathsf{one}}; \\
&\quad \texttt{(* load stack pointer *)} \\
&\quad \mathbf{Const}\ (\mathbf{stackl}\ c)\ r_{\mathsf{spp}}; \mathbf{Load}\ r_{\mathsf{spp}}\ r_{\mathsf{sp}}; \\
&\quad \texttt{(* push return address *)} \\
&\quad \mathbf{Add}\ r_{\mathsf{sp}}\ r_{\mathsf{one}}\ r_{\mathsf{sp}}; \mathbf{Store}\ r_{\mathsf{sp}}\ r_{\mathsf{a}}; \mathscr{A}(c, \mathit{Icode})
\end{aligned}
$$

where $\mathscr{A}(c, \mathit{Icode})$ is an auxiliary, recursively defined function having $\mathscr{A}(c, \texttt{[]}) = \texttt{[]}$ as a base case. As shown in the code snippet, the first instructions initialize the registers for them to match with the invariant we just explained informally.

Compilation is most interesting for calls and return instructions, which we present in figure 10. We

$$
\begin{aligned}
&\mathscr{A}(c, \mathbf{Call}\ c'\ m; \mathit{Icode}) = \\
&\quad \texttt{(* pop call argument and object *)} \\
&\quad \mathbf{Load}\ r_{\mathsf{sp}}\ r_{\mathsf{aux2}}; \mathbf{Sub}\ r_{\mathsf{sp}}\ r_{\mathsf{one}}\ r_{\mathsf{sp}}; \mathbf{Load}\ r_{\mathsf{sp}}\ r_{\mathsf{aux1}}; \\
&\quad \texttt{(* push current object and argument *)} \\
&\quad \mathbf{Store}\ r_{\mathsf{sp}}\ r_{\mathsf{tgt}}; \mathbf{Add}\ r_{\mathsf{sp}}\ r_{\mathsf{one}}\ r_{\mathsf{sp}}; \mathbf{Store}\ r_{\mathsf{sp}}\ r_{\mathsf{arg}}; \\
&\quad \texttt{(* save stack pointer *)} \\
&\quad \mathbf{Store}\ r_{\mathsf{spp}}\ r_{\mathsf{sp}}; \\
&\quad \texttt{(* set call object and argument *)} \\
&\quad \mathbf{Mov}\ r_{\mathsf{aux1}}\ r_{\mathsf{tgt}}; \mathbf{Mov}\ r_{\mathsf{aux2}}\ r_{\mathsf{arg}}; \\
&\quad \texttt{(* perform call *)} \\
&\quad \mathbf{Const}\ (\mathbf{methl}\ c'\ m)\ r_{\mathsf{aux3}}; \mathbf{Jal}\ r_{\mathsf{aux3}}; \\
&\quad \texttt{(* reinitialize environment *)} \\
&\quad \mathbf{Const}\ 1\ r_{\mathsf{one}}; \mathbf{Const}\ (\mathbf{stackl}\ c)\ r_{\mathsf{spp}}; \mathbf{Load}\ r_{\mathsf{spp}}\ r_{\mathsf{sp}}; \\
&\quad \texttt{(* restore current object and argument *)} \\
&\quad \mathbf{Load}\ r_{\mathsf{sp}}\ r_{\mathsf{arg}}; \mathbf{Sub}\ r_{\mathsf{sp}}\ r_{\mathsf{one}}\ r_{\mathsf{sp}}; \mathbf{Load}\ r_{\mathsf{sp}}\ r_{\mathsf{tgt}} \\
&\quad \texttt{(* push call result *)} \\
&\quad \mathbf{Store}\ r_{\mathsf{sp}}\ r_{\mathsf{ret}}; \mathscr{A}(c, \mathit{Icode})
\end{aligned}
$$

$$
\begin{aligned}
&\mathscr{A}(c, \mathbf{Ret}; \mathit{Icode}) = \\
&\quad \texttt{(* pop return value *)} \\
&\quad \mathbf{Load}\ r_{\mathsf{sp}}\ r_{\mathsf{ret}}; \mathbf{Sub}\ r_{\mathsf{sp}}\ r_{\mathsf{one}}\ r_{\mathsf{sp}}; \\
&\quad \texttt{(* pop return address *)} \\
&\quad \mathbf{Load}\ r_{\mathsf{sp}}\ r_{\mathsf{a}}; \mathbf{Sub}\ r_{\mathsf{sp}}\ r_{\mathsf{one}}\ r_{\mathsf{sp}}; \\
&\quad \texttt{(* save stack pointer *)} \\
&\quad \mathbf{Store}\ r_{\mathsf{spp}}\ r_{\mathsf{sp}}; \\
&\quad \texttt{(* perform return *)} \\
&\quad \mathbf{Jump}\ r_{\mathsf{a}}; \mathscr{A}(c, \mathit{Icode})
\end{aligned}
$$

Figure 10: Compilation of communication-related instructions of the intermediate machine

$$\mathscr{A}(c, \mathbf{This}; Icode) = \texttt{(* push object *)}$$
$$\quad \mathbf{Add}\, r_{\mathsf{sp}}\, r_{\mathsf{one}}\, r_{\mathsf{sp}}; \mathbf{Store}\, r_{\mathsf{sp}}\, r_{\mathsf{tgt}}; \mathscr{A}(c, Icode)$$

$$\mathscr{A}(c, \mathbf{Arg}; Icode) = \texttt{(* push argument *)}$$
$$\quad \mathbf{Add}\, r_{\mathsf{sp}}\, r_{\mathsf{one}}\, r_{\mathsf{sp}}; \mathbf{Store}\, r_{\mathsf{sp}}\, r_{\mathsf{arg}}; \mathscr{A}(c, Icode)$$

$$\mathscr{A}(c, \mathbf{Ref}\ o; Icode) = \texttt{(* push object o *)}$$
$$\quad \mathbf{Const}\ (\mathbf{objl}\ o)\, r_{\mathsf{aux1}};$$
$$\quad \mathbf{Add}\, r_{\mathsf{sp}}\, r_{\mathsf{one}}\, r_{\mathsf{sp}}; \mathbf{Store}\, r_{\mathsf{sp}}\, r_{\mathsf{aux1}}; \mathscr{A}(c, Icode)$$

$$\mathscr{A}(c, \mathbf{Drop}; Icode) = \mathbf{Sub}\, r_{\mathsf{sp}}\, r_{\mathsf{one}}\, r_{\mathsf{sp}}; \mathscr{A}(c, Icode)$$

$$\mathscr{A}(c, \mathbf{Sel}\, f; Icode) =$$
$$\quad \mathbf{Const}\ (f-1)\, r_{\mathsf{aux2}};$$
$$\quad \texttt{(* pop object to select from *)}$$
$$\quad \mathbf{Load}\, r_{\mathsf{sp}}\, r_{\mathsf{aux1}}; \mathbf{Add}\, r_{\mathsf{aux1}}\, r_{\mathsf{aux2}}\, r_{\mathsf{aux1}};$$
$$\quad \texttt{(* load and push field value *)}$$
$$\quad \mathbf{Load}\, r_{\mathsf{aux1}}\, r_{\mathsf{aux1}}; \mathbf{Store}\, r_{\mathsf{sp}}\, r_{\mathsf{aux1}}; \mathscr{A}(c, Icode)$$

$$\mathscr{A}(c, \mathbf{Upd}\, f; Icode) =$$
$$\quad \mathbf{Const}\ (f-1)\, r_{\mathsf{aux2}};$$
$$\quad \texttt{(* pop new field value and object *)}$$
$$\quad \mathbf{Load}\, r_{\mathsf{sp}}\, r_{\mathsf{aux3}}; \mathbf{Sub}\, r_{\mathsf{sp}}\, r_{\mathsf{one}}\, r_{\mathsf{sp}}; \mathbf{Load}\, r_{\mathsf{sp}}\, r_{\mathsf{aux1}};$$
$$\quad \texttt{(* perform update on object *)}$$
$$\quad \mathbf{Add}\, r_{\mathsf{aux1}}\, r_{\mathsf{aux2}}\, r_{\mathsf{aux1}}; \mathbf{Store}\, r_{\mathsf{aux1}}\, r_{\mathsf{aux3}};$$
$$\quad \texttt{(* push new field value *)}$$
$$\quad \mathbf{Store}\, r_{\mathsf{sp}}\, r_{\mathsf{aux3}}; \mathscr{A}(c, Icode)$$

Figure 11: Compilation of stack-related instructions of the intermediate machine

$$\mathscr{A}(c, \mathbf{Nop}; Icode) = \mathbf{Nop}; \mathscr{A}(c, Icode)$$

$$\mathscr{A}(c, \mathbf{Halt}; Icode) = \mathbf{Halt}; \mathscr{A}(c, Icode)$$

$$\mathscr{A}(c, \mathbf{Skip}\ k; Iinstr_1; \ldots; Iinstr_k; Icode) =$$
$$\quad \mathbf{Bnz}\, r_{\mathsf{one}}\, \mathscr{L}(Iinstr_1; \ldots; Iinstr_k);$$
$$\quad \mathscr{A}(c, Iinstr_1; \ldots; Iinstr_k; Icode)$$

$$\mathscr{A}(c, \mathbf{Skeq}\ k; Iinstr_1; \ldots; Iinstr_k; Icode) =$$
$$\quad \texttt{(* pop and compare objects *)}$$
$$\quad \mathbf{Load}\, r_{\mathsf{sp}}\, r_{\mathsf{aux2}}; \mathbf{Sub}\, r_{\mathsf{sp}}\, r_{\mathsf{one}}\, r_{\mathsf{sp}}; \mathbf{Load}\, r_{\mathsf{sp}}\, r_{\mathsf{aux1}};$$
$$\quad \mathbf{Sub}\, r_{\mathsf{sp}}\, r_{\mathsf{one}}\, r_{\mathsf{sp}}; \mathbf{Binop}_{=} r_{\mathsf{aux1}} r_{\mathsf{aux2}} r_{\mathsf{aux1}};$$
$$\quad \texttt{(* branch according to result *)}$$
$$\quad \mathbf{Bnz}\, r_{\mathsf{aux1}}\, \mathscr{L}(Iinstr_1; \ldots; Iinstr_k);$$
$$\quad \mathscr{A}(c, Iinstr_1; \ldots; Iinstr_k; Icode)$$

where $\mathscr{L}(Iinstr_1; \ldots; Iinstr_k)$ is the length of the sequence of compiled instructions corresponding to instructions $Iinstr_1; \ldots; Iinstr_k$, defined as:

$$\begin{aligned}
\mathscr{L}(\texttt{[]}) &= 0 \\
\mathscr{L}(\mathbf{Drop}; Icode) = \mathscr{L}(\mathbf{Nop}; Icode) &= 1 + \mathscr{L}(Icode) \\
\mathscr{L}(\mathbf{Halt}; Icode) = \mathscr{L}(\mathbf{Skip}\ n; Icode) &= 1 + \mathscr{L}(Icode) \\
\mathscr{L}(\mathbf{This}; Icode) = \mathscr{L}(\mathbf{Arg}; Icode) &= 2 + \mathscr{L}(Icode) \\
\mathscr{L}(\mathbf{Ref}\ o; Icode) &= 3 + \mathscr{L}(Icode) \\
\mathscr{L}(\mathbf{Sel}\, f; Icode) &= 5 + \mathscr{L}(Icode) \\
\mathscr{L}(\mathbf{Ret}; Icode) = \mathscr{L}(\mathbf{Skeq}\ n; Icode) &= 6 + \mathscr{L}(Icode) \\
\mathscr{L}(\mathbf{Upd}\, f; Icode) &= 7 + \mathscr{L}(Icode) \\
\mathscr{L}(\mathbf{Call}\ c'\ m; Icode) &= 18 + \mathscr{L}(Icode)
\end{aligned}$$

Figure 12: Compilation of control-related instructions of the intermediate machine

also present the compilation of stack-manipulating instructions in figure 11 and that of control-related instructions in figure 12. In all these figures, inline comments are provided so that the interested reader can get a quick understanding of what is happening.

More standard compilers typically use a global call stack and compile code under the assumption that other components will not break invariants (such as r_{one} always holding value 1) nor mess with the call stack. In our case, however, other components may be controlled by an attacker, which is incompatible with such assumptions. As a consequence, we use local stacks not only for intermediate results, but also to spill registers r_{tgt} and r_{arg} before performing a call. After the call, we restore the values of registers r_{one}, r_{spp} and r_{sp} that could have been overwritten by the callee, then fill registers r_{tgt} and r_{arg} from the local stack.

There is a lot of room for improvement in terms of compiler efficiency; having a code optimization pass would be very interesting in the future.

6.2 Micro-Policy Protecting Abstractions

Here, we first present the abstractions that our source language provides with respect to low-level machine code, and give intuition about how we can protect them using a micro-policy. Then, we present our actual micro-policy and explain how it effectively implements this intuition. We end with a description of the low-level loader, which sets the initial memory tags for the program.

6.2.1 Enforcing Class Isolation via Compartmentalization

Abstraction Because in the source language fields are private, classes have no way to read or to write to the data of other classes. Moreover, classes cannot read or write code, which is fixed. Finally, the only way to communicate between classes is to perform a method call.

In machine code, however, **Load**, **Store** and **Jump** operations can target any address in memory, including those of other components. This interaction must be restricted to preserve the class isolation abstraction.

Protection Mechanism To enforce class isolation, memory cells and the program counter get tagged with a class name, which can be seen as a color. The code and data belonging to class c get tagged with color c.

Load and **Store** instructions are restricted to locations having the same color as the current instruction. Moreover, the rules will compare the next instruction's compartment to that of the current instruction: Switching compartments is only allowed for **Jump** and **Jal**; however, we need *more* protection on these instructions because switching compartments should be further restricted: this is what we now present.

6.2.2 Enforcing Method Call Discipline using Method Entry Points, Linear Return Capabilities, and Register Cleaning

Abstraction In the source language, callers and callees obey a strict call discipline: a caller performs a method call, which leads to the execution of the method body, and once this execution ends evaluation proceeds with the next operation of the caller. In machine code, though, the **Jal** and **Jump** instructions can target any address.

Moreover, in the high-level language callers and callees give no information to each other except for the arguments and the result of the call. In the low-level machine, however, registers may carry extra information about the caller or the callee and their intermediate computational results. This suggests a need for *register cleaning* upon calls and returns.

Protection Mechanism On our machine, calls are done via **Jal** instructions, which store the return address in the r_{a} register, and returns by executing a **Jump** to the value stored in r_{a}. The first goal here is to ensure that a caller can only performs calls at method entry points, and that a callee can only return to the return address it was given in register r_{a} on the corresponding call.

To this end, we extend the memory tags to allow tagging a memory location as a method entry point. Then, we monitor **Jal** instructions so that they can only target such locations. This is enough to protect method calls. For returns, however, the problem is not that simple. In contrast with calls, for which all method entry points are equally valid, only one return address is the good one at any given point in time; it is the one that was transferred to the callee in the r_{a} register by the corresponding call. More precisely, because there can be nested calls, there is exactly one valid return address for each call depth.

We reflect this by tracking the current call depth in the PC tag: it starts at zero and we increment it on **Jal** instructions and decrement it on **Jump** instructions that correspond to returns. With such tracking,

we can now, upon a **Jal** instruction, tag register r_{a} to mark its content as the only valid return address for the current call depth. It is however crucial to make sure that when we go back to call depth n, there isn't any return capability for call depth $n+1$ idling in the system. We do this by enforcing that the tag on the return address is never duplicated, which makes it a *linear return capability*. The return capability gets created upon **Jal**, moved from register to memory upon **Store** and from memory to register upon **Load**, moved between registers upon **Mov**, and is finally destroyed upon returning via **Jump**. When it gets destroyed, we can infer that the capability has disappeared from the system, since there was always only one.

Now that our first goal is met, we can think about the second one: ensuring that compiled components do not transmit more information upon call and return than they do within the high-level semantics. Upon return, the distrusted caller potentially receives more information than in the high-level: Uncleaned registers could hold information from the compiled callee. Upon calls, the distrusted callee similarly receives information through registers, but has also an extra information to use: the content of register r_{a}, which holds the return address. This content leaks the identity of the compiled caller to the distrusted callee, while there is no way for a callee to know the identity of its caller in the high-level. Fortunately, the content of register r_{a} is already a specifically tagged value, and we already prevent the use of linear return capabilities for any other means than returning through it or moving it around.

Let us now review the general purpose registers which could leak information about our compiled partial programs. r_{tgt} and r_{arg} could leak the current object and argument to the distrusted caller upon return, but this is fine: the caller was the one who set them before the call, so he already knows them. Upon call, these registers do not leak information either since, according to the compilation scheme, they are already overwritten with call parameters. r_{ret} could leak a previous result value of the compiled caller to the distrusted callee upon call: it has to be cleaned. Upon return, however, and according to the compilation scheme, this register is already overwritten with the return value for the call. r_{aux1}, r_{aux2}, r_{aux3} could leak intermediate results of the computation upon return and have to be cleaned accordingly. Upon call, however, following the compilation scheme, they are already overwritten with information that the distrusted callee either will get or already has. r_{spp} could leak the identity of the compiled caller to the distrusted callee upon call, since it contains a pointer to the caller's local stack's memory region: it has to be cleaned. In the case of a return however, the identity of the compiled callee is already known by the distrusted caller so no new information will be leaked. r_{sp} could leak information about the current state of the local stack, as well as the identity of the compiled caller, and should accordingly be cleaned in both call and return cases. r_{one} will be known by the distrusted caller or the distrusted callee to always hold value 1, and thus won't leak any information. r_{a} is already protected from being leaked: it is overwritten both at call and return time. In the case of a call, it is overwritten upon the execution of the **Jal** instruction to hold the new return address. In the case of a return, according to the compilation scheme, it will be overwritten before performing the return to hold the address to which we are returning. This description concerns the current unoptimizing compiler; for an optimizing compiler the situation would be a little different: more information could be leaked, and accordingly more registers would have to be cleaned.

A first solution would be to have the compiler produce register reset instructions **Const** $0\, r$ for every register r that could leak information, before any external call or return instruction. However, this would be very expensive. This is one of the reasons why we have made the following assumption about our target symbolic micro-policy machine: The tags of some fixed registers (here, r_{ret}, r_{spp} and r_{sp} upon **Jal**, and r_{aux1}, r_{aux2}, r_{aux3} and r_{sp} upon **Jump**) can be updated in our symbolic micro-policy rules. We are thus by assumption able to clean the registers that might leak information, by using a special tag to mark these registers as cleared when we execute a **Jump** or a **Jal** instruction.

6.2.3 Enforcing Type Safety Dynamically

Abstraction Finally, in the source language callees and callers expect arguments or return values that are well-typed with respect to method signatures: We only consider high-level components that are statically well-typed and thus have to comply with the type interface they declare. At the machine code level, however, compromised components are untyped and can feed arbitrary machine words to uncompromised ones, without any a priori typing restrictions.

Protection Mechanism We use micro-policy rules to ensure that method arguments and return

values always have the type mentioned in type signatures. Fortunately, our type system is simple enough that we can encode types (which are just class names) as tags. Hence, we can build upon the call discipline mechanism above and add the expected argument type to the entry point tag, and the expected return type to the linear return capability tag.

Our dynamic typing mechanism relies on the loader to initialize memory tags appropriately. This initial tagging will be presented in detail after the micro-policy itself: The main idea is that a register or memory location holding an object pointer will be tagged as such, together with the type of the object. This dynamic type information is moved around with the value when it is stored, loaded or copied. One key for making this possible is that the **Const** (**objl** o) r instructions which put an object reference in a register are *blessed* with the type of this object according to the type declared for o: executing a blessed **Const** instruction will put the corresponding type information on the target register's tag.

Remember that we assume that we can check the next instruction tag in micro-policy rules. With dynamic type information available, we can then do type checking by looking at the tags of registers r_{tgt} and r_{arg} upon call, and that of register r_{ret} upon return. Upon call, we will compare with type information from the next instruction tag, which is the tag for the method entry point. Upon return, we will compare with type information from the return capability's tag. For these checks to be possible, we use the following assumption we made about our target symbolic micro-policy machine: The tags of some fixed registers (here, r_{tgt} and r_{arg} upon **Jal** and r_{ret} upon **Jump**) can be checked in our symbolic micro-policy rules.

6.2.4 Micro-Policy in Detail

As presented in section 5.3, a symbolic micro-policy is the combination of a collection of symbolic tags and a transfer function. We first detail our tag syntax. Then, we give the rules of our symbolic micro-policy, which define its transfer function [9]. Finally, we explain how the loader initially tags program memory following the program's export declarations.

Symbolic Micro-Policy Tags Our collection of symbolic micro-policy tags is presented in figure 13.

The tag on the program counter register is a natural number n which represents the current call depth.

t_{pc} ::= program counter tag
  n

t_{mem} ::= memory tag
  (bt, c, et, vt)

t_{reg} ::= register tag
  vt

bt ::= blessed tag
  B c | **NB**

et ::= entry point tag
  EP $c_{\mathsf{a}} \to c_{\mathsf{r}}$ | **NEP**

vt ::= regular or cleared value tag
  rt | $\bot$

rt ::= regular value tag
  Ret $n\, c_{\mathsf{r}}$ | **O** c | **W**

Figure 13: Symbolic micro-policy tags syntax

Memory tags (bt, c, et, vt) combine the various information we need about memory cells: First, a memory cell belongs to a compartment c. Its bt tag can be either **B** c' for it to be blessed with type c' — which means that it is a **Const** instruction which puts an object of type c' in its target register — or **NB** when it shouldn't be blessed. Similarly, its et tag can be either **EP** $c_{\mathsf{a}} \to c_{\mathsf{r}}$ when it is the entry point of a method of type signature $c_{\mathsf{a}}(c_{\mathsf{r}})$, or **NEP** when it is not. Finally, the vt tag is for the content of the memory cell, which can be either: a cleared value, tagged $\bot$; a return capability for going back to call depth n by providing a return value of type c_{r}, tagged **Ret** $n\, c_{\mathsf{r}}$; an object pointer of type c', tagged **O** c'; or a regular word, tagged **W**.

Tags on registers are the same vt tags as the ones for the content of memory cells: The content of the register is also tagged as a cleared value, a return capability, an object pointer, or a regular word.

Micro-Policy Rules Our micro-policy is presented in figure 14, where we use the meta notation **clear** vt for clearing return capabilities: It is equal to vt, unless vt is a return capability, in which case it is equal to $\bot$.

This micro-policy combines all the informal intuition we previously gave in sections 6.2.1, 6.2.2 and 6.2.3 into one transfer function. A notable optimization is that we use the tag on the next instruction

$\mathbf{Nop} : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt')\} \implies \{t_{\mathsf{pc}}{}' = n\}$
$\mathbf{Const}\, i\, r_{\mathsf{d}} : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt')\} \implies \{t_{\mathsf{pc}}{}' = n, t_{r_{\mathsf{d}}}{}' = \mathbf{W}\}$
$\mathbf{Const}\, i\, r_{\mathsf{d}} : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{B}\, c', c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt')\} \implies \{t_{\mathsf{pc}}{}' = n, t_{r_{\mathsf{d}}}{}' = \mathbf{O}\, c'\}$
$\mathbf{Mov}\, r_{\mathsf{s}}\, r_{\mathsf{d}} : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt'), t_{r_{\mathsf{s}}} = vt, t_{r_{\mathsf{d}}} = vt'\}$
$\implies \{t_{\mathsf{pc}}{}' = n, t_{r_{\mathsf{d}}}{}' = vt, t_{r_{\mathsf{s}}}{}' = \mathbf{clear}\, vt\}$
$\mathbf{Binop}_{op}\, r_1\, r_2\, r_{\mathsf{d}} : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt'), t_{\mathsf{r}1} = \mathbf{W}, t_{\mathsf{r}2} = \mathbf{W}\}$
$\implies \{t_{\mathsf{pc}}{}' = n, t_{r_{\mathsf{d}}}{}' = \mathbf{W}\}$
$\mathbf{Load}\, r_{\mathsf{p}}\, r_{\mathsf{d}} : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt'), t_{r_{\mathsf{p}}} = \mathbf{W}, t_{\mathsf{mem}} = (bt, c, et'', vt), t_{r_{\mathsf{d}}} = vt'\}$
$\implies \{t_{\mathsf{pc}}{}' = n, t_{r_{\mathsf{d}}}{}' = vt, t_{\mathsf{mem}}{}' = (bt, c, et'', \mathbf{clear}\, vt)\}$
$\mathbf{Store}\, r_{\mathsf{p}}\, r_{\mathsf{s}} : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt'), t_{r_{\mathsf{p}}} = \mathbf{W}, t_{r_{\mathsf{s}}} = vt, t_{\mathsf{mem}} = (bt, c, et'', vt')\}$
$\implies \{t_{\mathsf{pc}}{}' = n, t_{\mathsf{mem}}{}' = (\mathbf{NB}, c, et'', vt), t_{r_{\mathsf{s}}}{}' = \mathbf{clear}\, vt\}$
$\mathbf{Jump}\, r : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt'), t_{\mathsf{r}} = \mathbf{W}\} \implies \{t_{\mathsf{pc}}{}' = n\}$
$\mathbf{Jump}\, r : \{t_{\mathsf{pc}} = n+1, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c', et', rt'), t_{\mathsf{r}} = \mathbf{Ret}\, n\, c, t_{r_{\mathsf{ret}}} = \mathbf{O}\, c\}$
$\implies \{t_{\mathsf{pc}}{}' = n, t_{\mathsf{r}}{}' = \bot, t_{r_{\mathsf{aux}_1}}{}' = \bot, t_{r_{\mathsf{aux}_2}}{}' = \bot, t_{r_{\mathsf{aux}_3}}{}' = \bot, t_{r_{\mathsf{sp}}}{}' = \bot\}$ when $c \neq c'$
$\mathbf{Jal}\, r : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt'), t_{\mathsf{r}} = \mathbf{W}, t_{r_{\mathsf{a}}} = vt'\} \implies \{t_{\mathsf{pc}}{}' = n, t_{r_{\mathsf{a}}}{}' = \mathbf{W}\}$
$\mathbf{Jal}\, r : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c', \mathbf{EP}\, c_1 \to c_2, rt'), t_{\mathsf{r}} = \mathbf{W}, t_{r_{\mathsf{a}}} = vt', t_{r_{\mathsf{arg}_1}} = \mathbf{O}\, c, t_{r_{\mathsf{arg}_2}} = \mathbf{O}\, c_1\}$
$\implies \{t_{\mathsf{pc}}{}' = n+1, t_{r_{\mathsf{a}}}{}' = \mathbf{Ret}\, n\, c_2, t_{r_{\mathsf{ret}}}{}' = \bot, t_{r_{\mathsf{spp}}}{}' = \bot, t_{r_{\mathsf{sp}}}{}' = \bot\}$ when $c \neq c'$
$\mathbf{Bnz}\, r\, i : \{t_{\mathsf{pc}} = n, t_{\mathsf{ci}} = (\mathbf{NB}, c, et, \mathbf{W}), t_{\mathsf{ni}} = (bt', c, et', rt'), t_{\mathsf{r}} = \mathbf{W}\} \implies \{t_{\mathsf{pc}}{}' = n\}$

Figure 14: The rules of our symbolic micro-policy

to distinguish internal from cross-compartment calls and returns. Indeed, we don't need to enforce method call discipline nor to perform type checking on internal calls and returns. This is a crucial move with respect to both transparency and efficiency: It means that low-level components don't have to comply with the source language's abstractions for their internal computations, and moreover this lax monitoring will result in a lower overhead on execution time for internal steps, since there will be less cache misses.

Loader Initializing Memory Tags The loader first performs simple static checks: it must make sure that (1) no import declaration is left (the program is complete); (2) there exists a **methl** $c\, m$ region for each method m of each exported class c; (3) there exists a **stackl** c region for each exported class c; (4) there exists an **objl** o region for each exported object o; (5) all memory regions have a matching counterpart in the export declarations.

If all these checks succeed the loader proceeds and tags all program memory, following the export declarations: Every memory region is tagged uniformly with the tag of its compartment; which is c for **methl** $c\, m$ and **stackl** c memory regions, and the exported type of o for **objl** o memory regions. The first memory cell of each method region **methl** $c\, m$ gets tagged as an entry point with the exported type signature for method m of class c, while all other memory cells in the program get tagged as not being entry points. All locations holding an encoded **Const** (**objl** o) r instruction are tagged as blessed instructions storing a c object pointer in register r (**B** c), where c is the exported type of object o. All locations that hold a symbolic pointer **objl** o are tagged as storing a pointer to an object of class c (**O** c), where c is the exported type of object o. This applies to cells in both object and stack memory regions. The other stack memory cells are tagged as cleared values $\bot$, and all the remaining memory cells as regular words **W**.

7 Related Work

7.1 Secure Compilation

Secure compilation has been the topic of many works [4, 19, 22, 32], but only recently has the problem of targeting machine code been considered [6,32]. Moreover, all of these works focus on protecting a program from its context, rather than protecting mutually distrustful components like we do.

Abadi and Plotkin [4] formalized address space layout randomization as a secure implementation scheme for private locations, with probabilistic guarantees: They expressed it as a fully-abstract compilation between a source language featuring public and private locations, and a lower-level language in which

memory addresses are natural numbers. Follow-ups were presented by Jagadeesan et al. [22] and Abadi, Planul and Plotkin [2,3], with extended programming languages.

Fournet et al. [19] constructed a fully-abstract compiler from a variant of ML [40] to JavaScript. Their protection scheme protects one generated script from its context. A key point is that the protected script must have first-starter privilege: The first script which gets executed can overwrite objects in the global namespace, on which other scripts may rely. Hence, this scheme can't protect mutually distrustful components from each other, since only one component can be the first to execute.

The closest work to ours is recent [6,31,32] and ongoing [31,34] work by Agten, Patrignani, Piessens, et al. They target a set of hardware extensions, which they call protected module architectures [20, 28, 38]. Micro-policies could be used to define a protected module architecture, but they offer finer-grained protection: In our work, this finer-grained protection allows us to manage linear return capabilities and perform dynamic type-checking. Micro-policies also allow us to support a stronger attacker model of dynamic corruption, by means of a secure compilation of mutually distrustful components. As we discovered recently [24, 34], Patrignani et al. are currently trying to extend there previous work to ensure a secure compilation of mutually distrustful components. Our hope is that this parallel work can lead to interesting comparison and exchange, because the mechanisms we use are different: We believe that exploiting the full power of micro-policies can provide stronger guarantees and better performance than using micro-policies just as an instance of protected module architectures.

7.2 Multi-Language Approaches

In contrast with previous fully-abstract compilers where a single source component gets protected from its context, we protect linked mutually distrustful low-level components from each other.

One benefit of this approach is that components need not share a common source language. While our current protection mechanism is still deeply connected to our source language, in principle each component could have been written in a specific source language and compiled using a specific compiler.

It is actually common in real-life that the final program comes from a mix of components that were all written in different languages. Giving semantics to multi-language programs and building verified compilers that can provide guarantees for them is a hot topic, studied in particular by Ahmed et al. [7, 35] and Ramananandro et al. [36].

8 Discussion and Future Work

In this section, we discuss the limitations of our work and the generality of our approach, as well as future work.

8.1 Finite Memory and Full Abstraction

While memory is infinite in our high-level language, memory is finite in any target low-level machine. Our symbolic micro-policy machine is no exception: memory regions have a fixed finite size. This means that memory exhaustion or exposing the size of regions can break full abstraction.

Let us first recall how memory regions are used in this work. Our compiler translates method bodies from high-level expressions to machine code: Each method gets a dedicated memory region in the process, to store its compiled code. This code manipulates a stack that is local to the method's compartment; and this stack also gets its own memory region. Finally, each object gets a dedicated memory region, storing low-level values for its fields.

The first problem is potential exhaustion of the local stack's memory: When the stack is full and the program tries to put a new value on top, the machine will stop. This already breaks compiler correctness: Executing the compiled code for (**this**; o) will for example first try to put r_{tgt} on top on the full stack and hence stop the machine, when the high-level expression would simply return o to the caller. Full abstraction, which typically relies on a compiler correctness lemma, is broken as well: The low-level attacker can now distinguish between method bodies (**this**; o) and o, even though they have the same behavior in the high-level. One workaround would be to add one more intermediate step in the compilation chain, where the symbolic machine would have infinite memory regions: Full abstraction would be preserved with respect to this machine, and weakened when we move to finite memory regions. This workaround is the one taken by CompCert [27], which until very recently [30] only formalized and proved something about infinite memory models. A better but probably more difficult to implement solution would be to keep the current finite-memory machine, but make it explicitly in the property (e.g. compiler

correctness, full abstraction) that in cases such as resource exhaustion all bets are off.

The second problem is that the size of compiled method bodies, as well as the number of private fields of compiled objects, exactly match the size of their dedicated memory regions. This does not cause problems with the current memory model in which region locations exist in isolation of other regions. In future work, switching to a more concrete view of memory could lead to the exposure of information to the attacker: If a program region happens to be surrounded by attacker memory regions, then the attacker could infer the size of the program region and hence get size information about a method body or an object. Because there is no similar information available in the high-level, this will likely break full abstraction. The concrete loader could mitigate this problem, for example by padding memory regions so that all implementations of the same component interface get the same size for each dedicated memory region. This would be, however, wasteful in practice. Alternatively, we could weaken the property to allow leaking this information. For instance we could weaken full abstraction to say that one can only replace a compartment with an equivalent one that has the same sizes when compiled. This would weaken the property quite a lot, but it would not waste memory. There could also be interesting compromises in terms of security vs practicality, in which we pad to fixed size blocks and we only leak the number of blocks.

These problems are not specific to our compiler. Fournet et al. [19] target JavaScript and view stack and heap memory exhaustion as a side channel of concrete JavaScript implementations: It is not modeled by the semantics they give to JavaScript. Similarly, the key to the full abstraction result of Patrignani et al. [32, 33] (the soundness and completeness of their trace semantics) is given under the assumption that there is no overflow of the program's stack [33]. Patrignani et al. [32] also pad the protected program so that for all implementations to use the same amount of memory.

8.2 Efficiency and Transparency

Our micro-policy constrains low-level programs so as to prevent them from taking potentially harmful actions. However, we should make sure 1) that this monitoring has reasonable impact on the performance of these programs; and 2) that these programs are not constrained too much, in particular that benign low-level components are not prevented from interacting with compiled components.

A first, good step in this direction is that we don't enforce method call discipline nor type safety on internal calls and returns, but only on cross-compartment calls and returns. This is a good idea for both efficiency and transparency: Checks are lighter, leading to better caching and thus better performance; and low-level programs are less constrained, while still being prevented from taking harmful actions.

However, the constraints we set may still be too restrictive: For example, we enforce an object-oriented view on function calls and on data, we limit the number of arguments a function can pass through registers, and we force the programs to comply with our type system. This suggests the need for *wrappers*. Since internal calls and returns are not heavily monitored, we can define methods that respect our constraints and internally call the non-compliant benign low-level code: This low-level code can then take its non-harmful, internal actions without constraints — hence with good performance — until it internally returns to the wrapper, which will appropriately convert the result of the call before returning to its caller.

8.3 Future Work

The first crucial next step is to finish the full abstraction proof. As we explain in section 2, however, full abstraction does not capture the exact notion of secure compilation we claim to provide. We will thus formalize a suitable characterization and prove it, hopefully reusing lemmas from the full abstraction proof. Afterwards, we will implement the compiler and conduct experiments that could confirm or deny our hopes regarding efficiency and transparency.

There are several ways to extend this work. The most obvious would be to support more features that are common in object-oriented languages, such as dynamic allocation, inheritance, packages or exceptions. Another way would be to move to functional languages, which provide different, interesting challenges. Taking as source language a lambda-calculus with references and simple modules, would be a first step in this direction, before moving to larger ML subsets.

Finally, the micro-policy use in this work was built progressively, out of distinct micro-policies which we designed somewhat independently. Composing micro-policies in a systematic and correct way, without breaking the guarantees of any of the composed policies, is still an open problem that would be very

interesting to study on its own.

8.4 Scaling to Real-World Languages

Our micro-policy seems to scale up easily to more complicated languages, except for dynamic type checking which will be trickier.

Sub-typing, which arises with inheritance, would bring the first new challenges in this respect. Our dynamic type checking mechanism moreover requires encoding types in tags: When we move to languages with richer type systems, we will have to explore in more detail the whole field of research on dynamic type and contract checking [18].

Compartmentalization could easily be extended to deal with public fields, by distinguishing memory locations that hold public field values from other locations.

Dynamic allocation seems possible and would be handled by monitor services, setting appropriate tags on the allocated memory region. However, such tag initialization is expensive for large memory regions in the current state of the PUMP, and could benefit from additional hardware acceleration [44].

Finally, functional languages bring interesting challenges that have little to do with the work presented in this document, such as closure protection and polymorphism. We plan to study these languages and discover how micro-policies can help in securely compiling them.

□

A Appendix

A.1 Encoding Usual Types

```
export obj decl tt : Unit
export class decl tt { }

obj tt : Unit { }
class Unit { }
```

Figure 15: Encoding the unit type

Here we give a flavor of what programming looks like with our source language by encoding some familiar types using our class mechanism: the unit type in figure 15, booleans in figure 16, and bounded natural numbers in figure 17. Encoding unbounded natural numbers would be possible with dynamic allocation, which is not part of our source language at the moment.

For better understanding, we use a syntax with strings for names which is easily mapped to our source language syntax. We present the three encoded types as distinct components, resulting in quite verbose programs: Linking them together in the high-level would result in one partial program with three classes and no import declarations.

A.2 Source Semantics

The semantics we propose for the source language is a small-step continuation-based semantics. It is particularly interesting to present this variant because it is very close to our intermediate machine and can help understanding how the source to intermediate compilation works.

After loading, source programs become a pair of a class table CT and an initial configuration Cfg. The syntax for configurations is presented on figure 18.

A configuration $(OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, e)$ can be thought of as a machine state: OT is the object table, from which we fetch field values and which gets updated when we perform field updates. CS is the call stack, on top of which we store the current environment upon call. o_{t} is the current object and o_{a} the current argument. e is the current expression to execute and K the current continuation, which defines what we should do with the result of evaluating e.

Configurations can be reduced: The rules for the reduction $CT \vdash Cfg \longrightarrow Cfg'$ are detailed in figure 19. The class table CT is on the left of the turnstile because it does not change throughout the computation.

The initial configuration $(OT, [\,], o_{\mathsf{t}}, o_{\mathsf{a}}, [\,], e)$ features the program's object table OT, an empty call stack, and an empty continuation. The current expression to execute, e, is the body of the main method of the program, executing with appropriately-typed current object o_{t} and argument o_{a}. Since object and class names are natural numbers, an example choice which we take in our formal study is to say that the main method is method 0 of class 0, and that it should initially be called with object 0 of type 0 as both current object and argument.

Our reduction is deterministic. A program terminates with result o_{r} when there is a possibly empty reduction sequence from its initial configuration to a final configuration $(OT', [\,], o_{\mathsf{t}}, o_{\mathsf{a}}, [\,], o_{\mathsf{r}})$ or $(OT', CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (\mathbf{exit}\,\square) :: K, o_{\mathsf{r}})$. The type system ensures that everywhere a **exit** e expression is encountered, expression e has the same type as the

expected return type for the main method. Hence, when a program terminates with a value, the value necessarily has this particular type.

```
import obj decl tt : Unit
import class decl Unit { }

export obj decl t, f : Bool
export class decl Bool {
  Bool not(Unit),
  Bool and(Bool),
  Bool or(Bool)
}

obj t : Bool { }
obj f : Bool { }
class Bool {
  Bool not(Unit) { this == t ? f : t }
  Bool and(Bool) { this == t ? arg : f }
  Bool or(Bool) { this == t ? t : arg }
}
```

Figure 16: Encoding booleans

```
export obj decl zero, one, two, three : BNat4
export class decl BNat4 {
  BNat4 add(BNat4),
  BNat4 mul(BNat4 arg)
}

obj zero  : BNat4 { zero, one }
obj one   : BNat4 { zero, two }
obj two   : BNat4 { one,  three }
obj three : BNat4 { two,  three }
class BNat4 {
  BNat4 pred, succ;

  BNat4 add(BNat4) {
    arg == zero ?
      this : this.succ.add(arg.pred)
  }

  BNat4 mul(BNat4) {
    arg == zero ?
      zero : this.mul(arg.pred).add(this)
  }
}
```

Figure 17: Encoding bounded natural numbers

Cfg ::= reduction configurations
$(OT, CS, o_{\mathtt{t}}, o_{\mathtt{a}}, K, e)$

CS ::= call stack
$[]$ $\mid (o_{\mathtt{t}}, o_{\mathtt{a}}, K) :: CS$

K ::= continuations
$[]$ $\mid E :: K$

E ::= flat evaluation contexts
$\Box.f \mid \Box.f := e' \mid o.f := \Box \mid \Box.m(e') \mid o.m(\Box)$
$\mid \Box == e' \ ? \ e'' : e''' \mid o == \Box \ ? \ e'' : e''' \mid \Box; e$
$\mid \mathbf{exit}\,\Box$

Figure 18: Configuration syntax for the source language

$\boxed{CT \vdash Cfg \longrightarrow Cfg'}$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, \mathbf{this}) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, o_{\mathsf{t}})}\ \text{THIS}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, \mathbf{arg}) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, o_{\mathsf{a}})}\ \text{ARG}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, e.f) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, \Box.f :: K, e)}\ \text{SEL_PUSH}$$

$$\frac{\begin{array}{l} OT(o_{\mathsf{t}}) = \mathbf{obj}\, c_{\mathsf{t}}\{o_{\mathsf{t}1}, \ldots, o_{\mathsf{t}k}\} \\ OT(o) = \mathbf{obj}\, c_{\mathsf{t}}\{o_1, \ldots, o_l\} \\ 1 \leq f \leq l \end{array}}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, \Box.f :: K, o) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, o_f)}\ \text{SEL_POP}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, e.f := e') \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (\Box.f := e') :: K, e)}\ \text{UPD_PUSH}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (\Box.f := e) :: K, o) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (o.f := \Box) :: K, e)}\ \text{UPD_SWITCH}$$

$$\frac{\begin{array}{l} OT(o_{\mathsf{t}}) = \mathbf{obj}\, c_{\mathsf{t}}\{o_{\mathsf{t}1}, \ldots, o_{\mathsf{t}k}\} \\ OT(o) = \mathbf{obj}\, c_{\mathsf{t}}\{o_1, \ldots, o_{f-1}, o_f, o'_j, \ldots, o'_l\} \\ OT' = OT[o \mapsto \mathbf{obj}\, c\{o_1, \ldots, o_{f-1}, o'', o'_j, \ldots, o'_l\}] \end{array}}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (o.f := \Box) :: K, o'') \longrightarrow (OT', CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, o'')}\ \text{UPD_POP}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, e.m(e')) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, \Box.m(e') :: K, e)}\ \text{CALL_PUSH}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, \Box.m(e) :: K, o) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, o.m(\Box) :: K, e)}\ \text{CALL_SWITCH}$$

$$\frac{\begin{array}{l} OT(o_{\mathsf{t}}') = \mathbf{obj}\, c_{\mathsf{t}}'\{o_{\mathsf{t}1}', \ldots, o_{\mathsf{t}k}'\} \\ CT(c_{\mathsf{t}}') = \mathbf{class}\, \{c_1, \ldots, c_i; M_1, \ldots, M_j\} \\ 1 \leq m \leq j \\ M_m = c_{\mathsf{r}}(c_{\mathsf{a}})\{e\} \end{array}}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, o_{\mathsf{t}}'.m(\Box) :: K, o_{\mathsf{a}}') \longrightarrow (OT, (o_{\mathsf{t}}, o_{\mathsf{a}}, K) :: CS, o_{\mathsf{t}}', o_{\mathsf{a}}', K, e)}\ \text{CALL_POP}$$

$$\frac{}{CT \vdash (OT, (o_{\mathsf{t}}', o_{\mathsf{a}}', K) :: CS, o_{\mathsf{t}}, o_{\mathsf{a}}, [\,], o_{\mathsf{r}}) \longrightarrow (OT, CS, o_{\mathsf{t}}', o_{\mathsf{a}}', K, o_{\mathsf{r}})}\ \text{RETURN}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, e_1 == e_2\ ?\ e_3 : e_4) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (\Box == e_2\ ?\ e_3 : e_4) :: K, e_1)}\ \text{TEST_PUSH}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (\Box == e_2\ ?\ e_3 : e_4) :: K, o_1) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (o_1 == \Box\ ?\ e_3 : e_4) :: K, e_2)}\ \text{TEST_SWITCH}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (o_1 == \Box\ ?\ e_3 : e_4) :: K, o_1) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, e_3)}\ \text{TEST_POP_EQ}$$

$$\frac{o_1 \neq o_2}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (o_1 == \Box\ ?\ e_3 : e_4) :: K, o_2) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, e_4)}\ \text{TEST_POP_NEQ}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, (e; e')) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (\Box; e') :: K, e)}\ \text{SEQ_PUSH}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (\Box; e) :: K, o) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, e)}\ \text{SEQ_POP}$$

$$\frac{}{CT \vdash (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, K, \mathbf{exit}\ e) \longrightarrow (OT, CS, o_{\mathsf{t}}, o_{\mathsf{a}}, (\mathbf{exit}\ \Box) :: K, e)}\ \text{EXIT_PUSH}$$

Figure 19: Continuation-based semantics for the source language

References

[1] M. Abadi. Protection in programming-language translations. Research Report 154, SRC, 1998.

[2] M. Abadi and J. Planul. On layout randomization for arrays and functions. In D. A. Basin and J. C. Mitchell, editors, *Principles of Security and Trust - Second International Conference, POST 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings*. 2013.

[3] M. Abadi, J. Planul, and G. D. Plotkin. Layout randomization and nondeterminism. *Electr. Notes Theor. Comput. Sci.*, 298:29–50, 2013.

[4] M. Abadi and G. D. Plotkin. On protection by layout randomization. *ACM Trans. Inf. Syst. Secur.*, 15(2):8, 2012.

[5] P. Agten, B. Jacobs, and F. Piessens. Sound modular verification of C code executing in an unverified context. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, 2015.

[6] P. Agten, R. Strackx, B. Jacobs, and F. Piessens. Secure compilation to modern processors. In S. Chong, editor, *25th IEEE Computer Security Foundations Symposium, CSF 2012, Cambridge, MA, USA, June 25-27, 2012*. 2012.

[7] A. Ahmed. Verified compilers for a multi-language world. In T. Ball, R. Bodík, S. Krishnamurthi, B. S. Lerner, and G. Morrisett, editors, *1st Summit on Advances in Programming Languages, SNAPL 2015, May 3-6, 2015, Asilomar, California, USA*. 2015.

[8] A. Azevedo de Amorim, N. Collins, A. DeHon, D. Demange, C. Hriţcu, D. Pichardie, B. C. Pierce, R. Pollack, and A. Tolmach. A verified information-flow architecture. *POPL*. 2014.

[9] A. Azevedo de Amorim, M. Dénès, N. Giannarakis, C. Hriţcu, B. C. Pierce, A. Spector-Zabusky, and A. Tolmach. Micro-policies: Formally verified, tag-based security monitors. In *36th IEEE Symposium on Security and Privacy (Oakland S&P)*. 2015.

[10] G. M. Bierman and M. J. Parkinson. Effects and effect inference for a core java calculus. *Electr. Notes Theor. Comput. Sci.*, 82(7):82–107, 2003.

[11] N. P. Carter, S. W. Keckler, and W. J. Dally. Hardware support for fast capability-based addressing. In F. Baskett and D. W. Clark, editors, *ASPLOS-VI Proceedings - Sixth International Conference on Architectural Support for Programming Languages and Operating Systems, San Jose, California, USA, October 4-7, 1994*. 1994.

[12] T. Dai, S. Sathyanarayan, R. H. C. Yap, and Z. Liang. Detecting and preventing activex api-misuse vulnerabilities in internet explorer. In T. W. Chim and T. H. Yuen, editors, *Information and Communications Security - 14th International Conference, ICICS 2012, Hong Kong, China, October 29-31, 2012. Proceedings*. 2012.

[13] A. A. de Amorim and B. C. Pierce. Alpha is for address! the essence of memory safety. Draft, 2015.

[14] U. Dhawan, C. Hriţcu, R. Rubin, N. Vasilakis, S. Chiricescu, J. M. Smith, T. F. Knight, Jr., B. C. Pierce, and A. DeHon. Architectural support for software-defined metadata processing. *ASPLOS*, 2015.

[15] Ú. Erlingsson. Low-level software security: Attacks and defenses. In *Foundations of Security Analysis and Design IV, FOSAD 2006/2007 Tutorial Lectures*, 2007.

[16] Ú. Erlingsson, M. Abadi, M. Vrable, M. Budiu, and G. C. Necula. XFI: Software guards for system address spaces. *OSDI*. 2006.

[17] I. Evans, S. Fingeret, J. Gonzalez, U. Otgonbaatar, T. Tang, H. Shrobe, S. Sidiroglou-Douskos, M. Rinard, and H. Okhravi. Missing the point(er): On the effectiveness of code pointer integrity. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, 2015.

[18] R. B. Findler and M. Felleisen. Contracts for higher-order functions. In *Proceedings of the 7th International Conference on Functional Programming*. 2002.

[19] C. Fournet, N. Swamy, J. Chen, P. Dagand, P. Strub, and B. Livshits. Fully abstract compilation to javascript. In R. Giacobazzi and

R. Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*. 2013.

[20] M. Hoekstra, R. Lal, P. Pappachan, V. Phegade, and J. del Cuvillo. Using innovative instructions to create trustworthy software solutions. In Lee and Shi [26].

[21] A. Igarashi, B. C. Pierce, and P. Wadler. Featherweight java: a minimal core calculus for java and GJ. *ACM Trans. Program. Lang. Syst.*, 23(3):396–450, 2001.

[22] R. Jagadeesan, C. Pitcher, J. Rathke, and J. Riely. Local memory via layout randomization. In *Proceedings of the 24th IEEE Computer Security Foundations Symposium, CSF 2011, Cernay-la-Ville, France, 27-29 June, 2011*. 2011.

[23] A. Jeffrey and J. Rathke. Java jr: Fully abstract trace semantics for a core java language. In S. Sagiv, editor, *Programming Languages and Systems, 14th European Symposium on Programming,ESOP 2005, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings*. 2005.

[24] Y. Juglaret and C. Hriţcu. Secure compilation using micro-policies (extended abstract). Workshop on Foundations of Computer Security, 2015.

[25] A. Kennedy. Securing the .net programming model. *Theor. Comput. Sci.*, 364(3):311–317, 2006.

[26] R. B. Lee and W. Shi, editors. *HASP 2013, The Second Workshop on Hardware and Architectural Support for Security and Privacy, Tel-Aviv, Israel, June 23-24, 2013*. ACM, 2013.

[27] X. Leroy. Formal verification of a realistic compiler. *CACM*, 52(7):107–115, 2009.

[28] F. McKeen, I. Alexandrovich, A. Berenzon, C. V. Rozas, H. Shafi, V. Shanbhogue, and U. R. Savagaonkar. Innovative instructions and software model for isolated execution. In Lee and Shi [26].

[29] G. Morrisett, G. Tan, J. Tassarotti, J.-B. Tristan, and E. Gan. RockSalt: better, faster, stronger SFI for the x86. *PLDI*. 2012.

[30] E. Mullen, Z. Tatlock, and D. Grossman. Peek: A formally verified peephole optimization framework for x86. CoqPL Workshop, 2015.

[31] M. Patrignani. *The Tome of Secure Compilation: Fully Abstract Compilation to Protected Modules Architectures*. PhD thesis, KU Leuven, Leuven, Belgium, 2015.

[32] M. Patrignani, P. Agten, R. Strackx, B. Jacobs, D. Clarke, and F. Piessens. Secure compilation to protected module architectures. *ACM Transactions on Programming Languages and Systems*, 2015.

[33] M. Patrignani and D. Clarke. Fully abstract trace semantics for protected module architectures. *Computer Languages, Systems & Structures*, 42(0):22 – 45, 2015. Special issue on the Programming Languages track at the 29th {ACM} Symposium on Applied Computing.

[34] M. Patrignani, D. Devriese, and F. Piessens. Multi-module fully abstract compilation (extended abstract). Workshop on Foundations of Computer Security, 2015.

[35] J. T. Perconti and A. Ahmed. Verifying an open compiler using multi-language semantics. In Z. Shao, editor, *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*. 2014.

[36] T. Ramananandro, Z. Shao, S. Weng, J. Koenig, and Y. Fu. A compositional semantics for verified separate compilation and linking. In X. Leroy and A. Tiu, editors, *Proceedings of the 2015 Conference on Certified Programs and Proofs, CPP 2015, Mumbai, India, January 15-17, 2015*. 2015.

[37] K. Z. Snow, F. Monrose, L. Davi, A. Dmitrienko, C. Liebchen, and A. Sadeghi. Just-in-time code reuse: On the effectiveness of fine-grained address space layout randomization. In *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*, 2013.

[38] R. Strackx and F. Piessens. Fides: selectively hardening software application components against kernel-level or process-level malware. In T. Yu, G. Danezis, and V. D. Gligor, editors, *the*

ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012. 2012.

[39] G. T. Sullivan, S. Chiricescu, A. DeHon, D. Demange, S. Iyer, A. Kliger, G. Morrisett, B. C. Pierce, H. Reubenstein, J. M. Smith, A. Thomas, J. Tov, C. M. White, and D. Wittenberg. SAFE: A clean-slate architecture for secure systems. In *Proceedings of the IEEE International Conference on Technologies for Homeland Security*, 2013.

[40] N. Swamy, J. Chen, C. Fournet, P. Strub, K. Bhargavan, and J. Yang. Secure distributed programming with value-dependent types. *J. Funct. Program.*, 23(4):402–451, 2013.

[41] L. Szekeres, M. Payer, T. Wei, and D. Song. SoK: Eternal war in memory. *IEEE S&P*. 2013.

[42] R. N. M. Watson, P. G. Neumann, J. Woodruff, J. Anderson, R. Anderson, N. Dave, B. Laurie, S. W. Moore, S. J. Murdoch, P. Paeps, M. Roe, and H. Saidi. CHERI: a research platform deconflating hardware virtualization and protection. In *Proc. RESoLVE*, 2012.

[43] R. N. M. Watson, J. Woodruff, P. G. Neumann, S. W. Moore, J. Anderson, D. Chisnall, N. H. Dave, B. Davis, K. Gudka, B. Laurie, S. J. Murdoch, R. Norton, M. Roe, S. Son, and M. Vadera. CHERI: A hybrid capability-system architecture for scalable software compartmentalization. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, 2015.

[44] E. Witchel, J. Cates, and K. Asanović. Mondrian memory protection. *ASPLOS*. 2002.

[45] B. Yee, D. Sehr, G. Dardyk, J. B. Chen, R. Muth, T. Ormandy, S. Okasaka, N. Narula, and N. Fullagar. Native client: a sandbox for portable, untrusted x86 native code. *Commun. ACM*, 53(1):91–99, 2010.